AD-A074 625 MASSACHUSETTS INST OF TECH CAMBRIDGE LAB FOR COMPUTE--ETC F/G 5/7
ROLES, CO-DESCRIPTORS, AND THE FORMAL REPRESENTATION OF QUANTIF--ETC(U)
SEP 79 W A MARTIN N00014-75-C-0661
UNCLASSIFIED MIT/LCS/TM-139 NL

| OF |
AD
A074625

END
DATE
FILMED
10-79
DDC

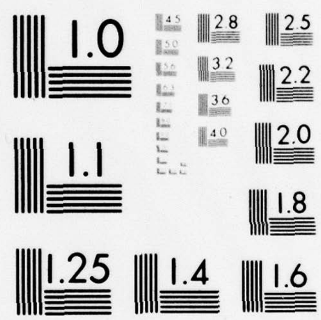
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0

LABORATORY FOR COMPUTER SCIENCE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

MIT/LCS/TM-139

ROLES, CO-DESCRIPTORS, AND THE FORMAL REPRESENTATION OF QUANTIFIED ENGLISH EXPRESSIONS

William A. Martin

September 1979

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Office of Naval Research under Contract No. N00014-75-C-0661

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER (14) MIT/LCS/TM-139	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) (6) Roles, Co-Descriptors, and the Formal Representation of Quantified English Expressions		5. TYPE OF REPORT & PERIOD COVERED
		6. PERFORMING ORG. REPORT NUMBER MIT/LCS/TM-139
7. AUTHOR(s) (10) William A. Martin		8. CONTRACT OR GRANT NUMBER(s) (15) N00014-75-C-0661
9. PERFORMING ORGANIZATION NAME AND ADDRESS MIT/Laboratory for Computer Science 545 Technology Square Cambridge, MA 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS ARPA/Department of Defense 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE (11) September 1979
		13. NUMBER OF PAGES 80
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) ONR/Department of the Navy Information Systems Program Arlington, VA 22217		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

(12) 81

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Knowledge Representation
Quantifiers
Semantic Nets

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

In representing the semantics of English sentences it is traditional to distinguish logical form from semantic content. The logical form is represented by some sort of predicate calculus or lambda calculus. In computational linguistics this predicate calculus or lambda calculus notation is usually carried over:

a) directly;

b) by replacing parenthesized scope with 'contexts'; (Hendrix 1978) <1> → over

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

409648

VP

20. c) by replacing universal quantification with dynamically scoped iteration procedures (Woods 1977).

This paper proposes another possibility for representing logical form. It is based on five main ideas:

a) The use of roles in a semantic net;
b) The referential / attribute distinction;
c) The distributive / collective distinction;
d) The use of predicates taking sets, kinds, or prototypical individuals as arguments; and
e) The use of two levels of representation for quantified expressions.

Roles, Co-Descriptors, and the Formal Representation of Quantified English Expressions

by

William A. Martin

Laboratory for Computer Science
Massachusetts Institute of Technology

Acknowledgement

Ken Church, Lowell Hawkinson, Mitchell Marcus, Peter Szolovits, and Lucia Vania read an earlier version of this manuscript and made many helpful comments. Ellen Lewis did an excellent job of preparing the manuscript and figures. This research was supported by the Defense Advance Research Projects Agency and monitored by the Office of Naval Research under contract no. N00014-75-C-0661.

Key Words: Knowledge Representation
Quantifiers
Semantic Nets

Contents

1. Introduction

In representing the semantics of English sentences it is traditional to distinguish logical form from semantic content. The logical form is represented by some sort of predicate calculus or lambda calculus. In computational linguistics this predicate calculus or lambda calculus notation is usually carried over

a) directly

b) by replacing parenthesized scope with 'contexts' (Hendrix 1978) <1>

c) by replacing universal quantification with dynamically scoped iteration procedures (Woods 1977)

This paper proposes another possibility for representing logical form. It is based on five main ideas.

a) The use of roles in a semantic net.

b) The referential / attributive distinction.

c) The distributive / collective distinction.

d) The use of predicates taking sets, kinds, or prototypical individuals as arguments.

e) The use of two levels of representation for quantified expressions.

The proposed scheme uses more mechanisms than older ones, but I will argue that it captures subtle cases in a more natural way, and that it is a superior representation for computational use.

2. Roles

The concept of a role has been introduced frequently by designers of knowledge representation languages taking the form of a semantic network (e.g. Martin 1979, Fahlman 1979, Brachman 1978). The notion of role is illustrated by the simplified semantic network in Figure 1.

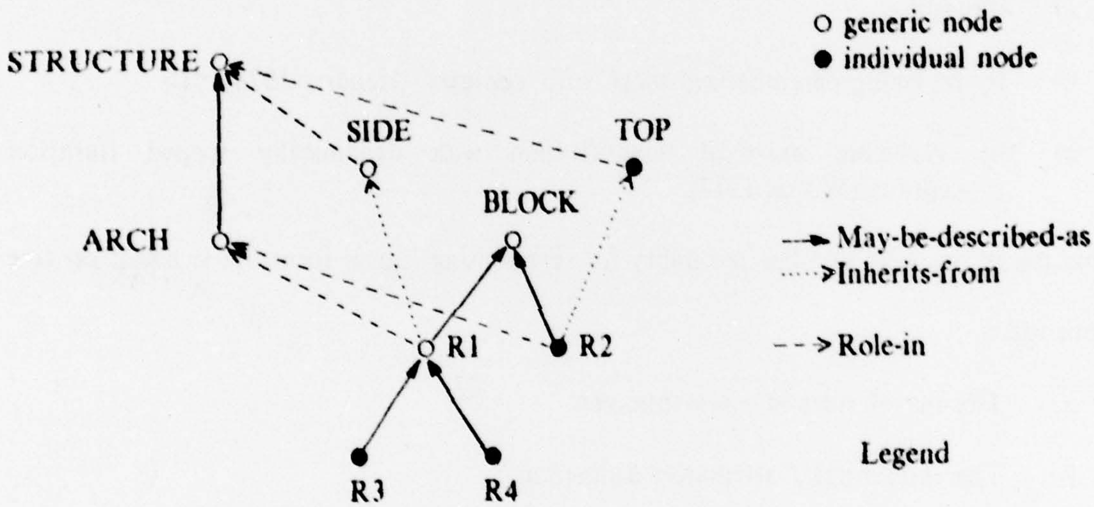

Figure 1
Simplified Structural Description of an Arch

This network has two kinds of nodes and three kinds of links. Hollow circles represent generic nodes, solid circles represent individual nodes. A generic node is the locus of a generic description of some type of entity and serves the function of what others have termed a frame, unit, or concept. An individual node differs from a generic node only in that while there may be any number of entities of the type described by a generic node, there is at most one entity which meets the description of an individual node.

We say "at most one", because the network may contain information that an individual meeting that node's description doesn't exist.

The description of a node is inferred from the rest of the network by tracing out the links. Nodes in the structural description of a given node, e.g. R1 in ARCH, are linked to the given node by role-in links. These nodes are termed roles in the description of the given node. As illustrated in Figure 1, a role may be either generic or individual.

One node, e.g. ARCH, may be described as another node, e.g. STRUCTURE, if for any individual which would meet the description of the first node we are willing to describe it by the second node. We are willing to describe a SQUARE as a RECTANGLE because a SQUARE is defined as a RECTANGLE with four equal sides. The relation holds by definition. SQUARE and RECTANGLE are said to be related on the basis of analytic knowledge. More commonly our willingness to describe one node as another is based on assumptions about the world, e.g. that COW may be described as DOMESTIC-ANIMAL requires us to know that the world doesn't contain wild cows (at least not enough of them to affect our description of the generic cow). Such knowledge of the world is termed synthetic knowledge. Our willingness to describe one node as another is generally based on a combination of analytic and synthetic knowledge. The analytic/synthetic distinction is not a basic distinction in our network.

The may-be-described-as relation models the subset relation. Instead of saying that a square may be described as a rectangle, we could have introduced notation to say that the set of all squares is a subset of the set of all rectangles. However, we prefer not to introduce, as a fundamental construct, the notion of the set of all entities meeting a certain description. We will work instead with generic nodes (as intensions) and individual nodes (as extensions). Quine (1970) points out that a system like ours still allows reasoning of the form done with the Boolean operations of set theory, without introducing Russelll's paradox. It also lets us side-step problems arising in the discussion of non-existent entities.

Using these conventions we may interpret Figure 1 as saying that an ARCH may be described as a STRUCTURE. Its structural description consists of an individual node, R2, which may be described as a BLOCK and inherits the description of a TOP, a generic, R1, which may be described as a BLOCK and inherits the description of a SIDE, and two individual nodes, R3 and R4, which may be described as the generic R1.

We make the convention that since R3 and R4 may be described as R1, and R1 is a role in ARCH, then R3 and R4 are also roles in ARCH.

Note that the roles of a node correspond to the slots of a frame in a theory like Minsky's (1975), as implemented in FRL by Goldstein and Roberts (1977). However, there is a significant difference. Since slots are not frames themselves, slots and frames are not described by the same conventions. But since roles are nodes, they may be

described by applying the node description conventions recursively. The ideas in this paper could also be formulated representing frames with nodes and slots with links. <2>

2.1 Structure and Context

A second point to note is that this network does not deal with contextual problems by assuming a set of objects called contexts or partitions and then somehow stipulating for each node and link what contexts or partitions it is present in. Instead, the position is taken that roles actually allow a more natural modelling of the phenomena partitions were designed to model. For example, Figure 2 shows how Hendrix (1978) would represent the fact that legal persons can own physical objects. In this figure rectangles (including the square) represent partitions, ovals represent sets, and circles represent individuals. An individual or set is in a partition if its representing circle or oval is in the rectangle representing that partition. In this figure the individual, I, is an implication that any individual owning, X, has an individual agent, Y, which is a legal person, an individual object, Z, which is a physical object, an individual start-time, t_1, which is a time, and an individual end-time, t_2, which is also a time.

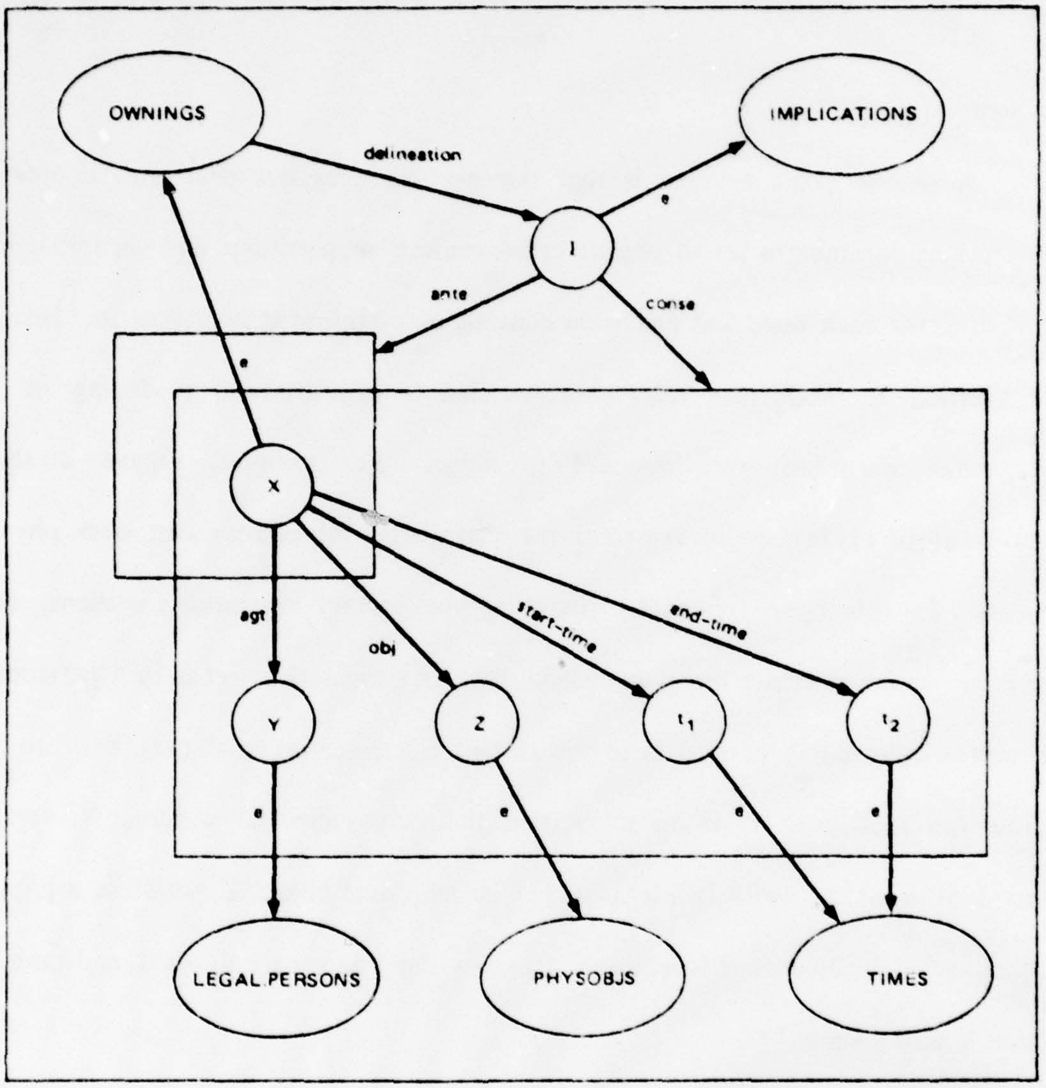

FIGURE 2 THE DELINEATION THEOREM OF OWNINGS

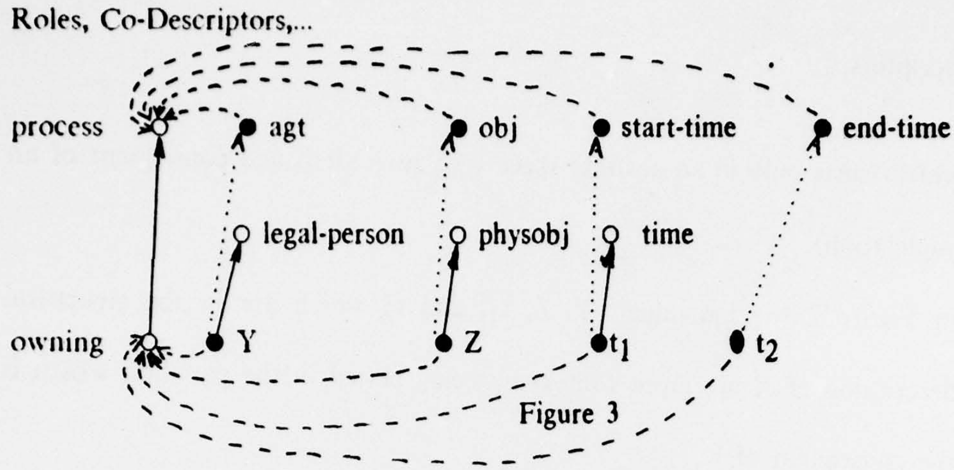

Figure 3

This same information is shown in Figure 3, using the notation presented here. The key differences between the representations in Figures 2 and 3 are:

a) i) Figure 2 shows the set of ownings. To indicate properties which must be true of every member of this set, an individual implication I is created. I is related to OWNINGS by the link, <u>delineation</u>. The antecedent of I is a partition containing an individual, X. By convention this individual, X, in the antecedent of I is the typical member of ownings. The individuals in the structural description of X are in the partition which is the consequent of I.

ii) In Figure 3, the node, owning, represents the typical member. As a generic representing the typical owning it does not need to be contextually constrained. Thus, it is not taken to be a role in any other node. Part of the complexity of Figure 2 is avoided by making the typical member a primitive notion (generic node) rather than constructing it as an individual

which exists only in an abstract space (the antecedent and consequent of an implication).

b) i) In Figure 2 the individuals Y, Z, t_1, and t_2 which are in the structural description of X are given context by being placed in the partition which is the consequent of I.

ii) In Figure 3 the individuals Y, Z, t_1, and t_2 which are in the structural description of X are given context by being made roles in X itself. Thus a typical member serves as the context of its own structural description.

c) In Figure 2, the individuals X, Z, t_1 and t_2 are related to X by links labeled with "case relation names" or "slot names". Thus, as mentioned above, a distinction is made between nodes and slots. By contrast in Figure 3, agt, obj, start-time, and end-time are themselves nodes and can be described just like any other node. It is perhaps worth pointing out that the logical organization of Figure 3 does not prevent memory from being physically organized so that given owning and agt one can retrieve node Y at least as quickly as it can be done with the organization of Figure 2.

In summary, Figure 2 takes as primary the notion of context and implements it by placing nodes in partitions. Figure 3 takes as primary the notion of structure and implements it by making a node a role in the structural description of another node. Probably all will agree that both structure and context are important notions for the

organization of information. The question is whether one or both of these notions should be reflected by the most primitive constructs of a representation language. By choosing just context, Hendrix is forced to go to a formal abstract notion - he picked implication in order to represent a structural description. Some authors have avoided this by taking both as equally primitive (Fahlman 1979, Hayes 1977).

Figure 4 shows how an entity can satisfy more than one individual node at the same time. This figure specifies that whatever is described as Bill is also described as the agent of kiss and the object of slap. This entity is the individual agent in the kiss at the same time that it is the individual object in the slap. (In this figure I have left out the links to generic nodes and the role-in links of Bill and Susie. I will continue to leave out links which are not the focus of discussion in order to simplify the figures.)

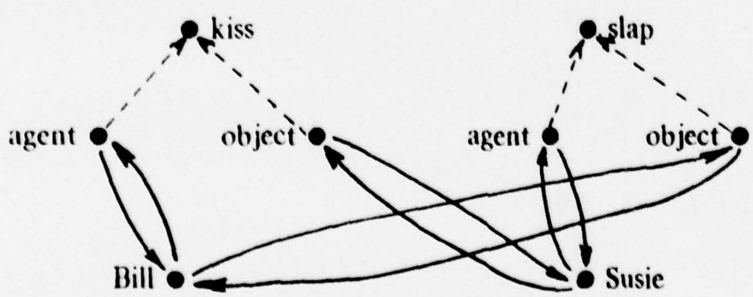

Figure 4
Bill kisses Susie and Susie slaps Bill

2.2 Inheritance

Fahlman (1979) has pointed out an interesting question which arises in the interpretation of structures like that shown in Figure 1, the structural description of an arch. If this figure is a description of the structure of the generic arch, then it should be possible to put this description into correspondence with the description of an individual arch, say ARCH-1. Figure 5 shows a fragment of Figure 1 extended to include the corresponding fragments of two individual arches, ARCH-1 and ARCH-2.

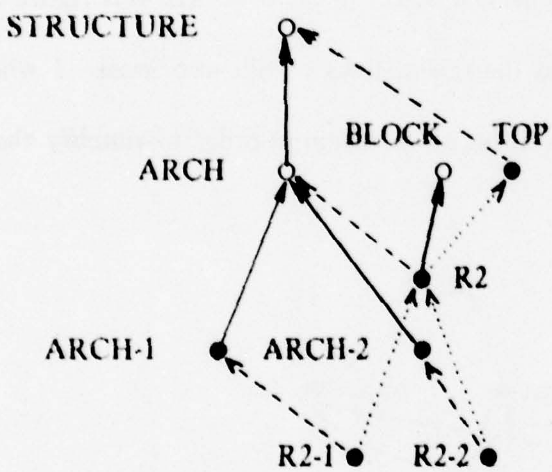

Figure 5

The dotted links from R2-1 to R2 and R2-2 to R2 specify inheritance. In this case R2-1 and R2-2 inherit all the description of R2 except the role-in link from R2 to ARCH. This role-in link is overridden by their own role-in links.

Inheritance links specify the inheritance of all description which is not explicitly overridden. May-be-described-as links indicate an alternative description. One might think that the inheritance links are not needed in Figure 5, that they could be replaced with may-be-described-as links. However, this leads to a problem pointed out by Fahlman.

Suppose we were to say that since R2-1 is the top of an individual arch, ARCH-1, it may be described as R2, the top of an arch. That is, when one node is described as another, its roles are described as the corresponding roles of the other. Doing this gets one into difficulty. To see this note that we may ask "which individual block in the generic arch is the top", and the answer is R2. R2 is taken as an individual node and clearly it represents a single one of the three blocks comprising an arch. But now suppose that nodes ARCH-1 and ARCH-2 describe arches in the real world. Then R2-1 and R2-2 describe the tops of these arches, and since R2-1 and R2-2 may be described as R2, R2 describes the tops of these arches. Thus, it would seem that R2 does not, in fact, describe a single individual. The reader will no doubt be able to articulate to his satisfaction what is going on in this example. We hope to capture this by expressing and invoking a rather far-reaching principle of interpretation.

Accordingly, we will stipulate:

> Every individual node, R, must be a role in some other node, N. Any entity, r, described by the node, R, is only taken to be an individual with respect to the entity, n, described by N. The only exception is that the real world is not in the structure of anything.

One effect of this stipulation is to make the role-in link an essential part of any description of an individual. Now by using inheritance links instead of may-be-described-as links from R2-1 to R2 and from R2-2 to R2, we can override the role-in link from R2 to ARCH and thus preserve R2 from describing these individuals.

The essential difference between inheritance and alternative description is that an inheriting node can override a criterial part of an inherited description, so that the inherited description no longer picks out the same individuals. By contrast, alternative description preserves the integrity of individuals. It lets one get a different description of the same individual.

2.3 Making Every Individual Node a Role

This section will touch briefly on the philosophical ramifications of making every individual node a role. In the Tarskian (1944) view of English semantics correspondence between descriptions and the world is described in terms of propositions. A proposition is true if it corresponds to the world and false otherwise. Propositions are either atomic or compound - formed by conjunction, etc. The truth or falsity of atomic propositions must be given. The truth or falsity of compound propositions is then determined from the truth or falsity of their components by applying rules of composition associated with their connectives.

In the theory proposed here, discussion of correspondence between descriptions

and the world is not limited to propositions. We can discuss correspondence for any individual, proposition or not <3>. However, the key difference is that since every description of an individual must include a description of the structure containing it, it doesn't make sense to ask for the truth or falsity of atomic propositions independently of the truth or falsity of the structures containing them. The view presented here shifts away from composition to a system which places more weight on decomposition. One starts with a node representing the world and structurally decomposes it in various ways. The parts of these decompositions are then further decomposed in turn to any desired depth. This produces a structural hierarchy of descriptions. This structural hierarchy is complemented by the abstraction hierarchy of descriptions produced by the may-be-described-as links.

The importance of this approach comes out if one believes that there are, in fact, individuals which have no description which does not depend criterially on their being part of a structure. This realization has caused some philosophers to deny that such individuals are entities at all. This view is put forward by Benacerraf (1965) in debating the existance of numbers:

> Therefore, numbers are not objects at all, because in giving the properties (that is, necessary and sufficient) of numbers you merely characterize an *abstract structure* -- and the distinction lies in the fact that the "elements" of the structure have no properties other than those relating them to other "elements" of the same structure. If we identify an abstract structure with a system of relations (in intension, of course, or else with the set of all relations in extension isomorphic to a given system of relations), we get arithmetic elaborating the properties of the "less-than" relation, or of all

> systems of objects (that is, *concrete* structures) exhibiting that abstract structure. That a system of objects exhibits the structure of the integers implies that the elements of that system have some properties not dependent on structure. It must be possible to individuate those objects independently of the role they play in that structure. But this is precisely what cannot be done with the numbers. To *be* the number 3 is no more or less than to be preceded by 2, 1, and possibly 0, and to be followed by 4, 5, and so forth. And to *be* the number 4 is no more and no less than to be preceded by 3, 2, 1, and possibly 0, and to be followed by .... *Any* object can *play the role of* 3; that is, any object can be the third element in some progression. What is peculiar to 3 is that it defines that role -- not by being a paradigm of any object which plays it, but by representing the relation that any third member of a progression bears to the rest of the progression.
>
> Arithemetic is therefore the science that elaborates the abstract structure that all progressions have in common merely in virtue of begin progressions. It is not a science concerned with particular objects -- the numbers. The search for which independently identifiable particular objects the numbers really are (sets? Julius Caesars?) is a misguided one.
>
> On this view many things that puzzled us in this paper seem to fall into place. Why so many interpretations of number theory are possible without any being uniquely singled out becomes obvious. There is no unique set of objects that are the numbers. Number theory is the elaboration of the properties of *all* structures of the order type of the numbers. The number words do not have single referents. Furthermore, the reason identification of numbers with objects works wholesale but fails utterly object by object is the fact that the theory is elaborating an abstract structure and not the properties of independent individuals, any one of which could be characterized without reference to its relations to the rest. Only when we are considering a particular sequence as being, not the numbers, but *of the structure of the numbers* does the question of which element is, or rather *corresponds to*, 3 begin to make any sense.

While I agree with Benacerraf's analysis of numbers, I am unwilling to therefore deny that numbers are entities. In a semantic network it is simplest to have nodes represent entities and one needs nodes for anything which will be described, including

numbers. <4> Rather than deny that numbers are entities I prefer to put all entities on the same footing as numbers - to require the description of any individual to describe the structure in which it is used. For some individuals, like numbers, the structure in which they are used will dominate their description. Other entities will have a complex structural description of their own.

3. Reference and Definite Descriptions

The above quote by Benacerraf deals with numbers in part because the philosophy of mathematics has historically been the main proving ground for formal representation systems. Our concern here, however, is with the logical structure of English sentences, and enough constructs have now been presented to permit a treatment of some well known phenomena. We begin with a discussion of definite descriptions.

3.1 Co-Referent Descriptions

Russell's (1905) analysis of a singular definite description required that there be a unique object satisfying the description in order for the expression to denote anything, and hence it fails to account for the successful reference of a noun phrase like "the clock" in

Did you wind the clock?

The problem is that "the clock" only picks out an individual clock given the context of

a particular use of the expression. Thus the expression picks out an individual only when paired with a description of the context of use. It is widely recognized that some method of describing this context has to be made available.

The general importance of context in forming descriptions of individuals may be seen in Strawson's essay, "Individuals" (1959). Strawson begins with a discussion relevant to the above difficulty with definite descriptions:

> When shall we say that a hearer knows what particular is being referred to by a speaker? Consider first the following case. A speaker tells a story which he claims to be factual. It begins, 'A man and a boy were standing by a fountain,' and it continues: 'The man had a drink.' Shall we say that the hearer knows which or what particular is being referred to by the subject expression in the second sentence? We might say so. For, of a certain range of two particulars the words 'the man' serve to distinguish the one being referred to, by means of a description which applies only to him. But though this is, in a weak sense, a case of identification, I shall call it only a story-relative or, for short, a relative identification. For it is identification only relative to a range of particulars (a range of two members) which is itself identified only as the range of particulars being talked about by the speaker. That is to say, the hearer, hearing the second sentence, knows which particular creature is being referred to of the two particular creatures being talked about by the speaker; but he does not, without this qualification, know what particular creature is being referred to. The identification is within a certain story told by a certain speaker. It is identification within his story but not identification within history.
>
> We need a requirement stringent enough to eliminate relative identification. The hearer, in the example, is able to place the particular referred to within the picture painted by the speaker. This means that in a sense he can place the particular in his own general picture of the world. For he can place the speaker, and hence the speaker's picture, in that general picture of his own. But he cannot place the figures, without the frame of the speaker's picture in his own general picture of the world. For this reason the full requirement of hearer's identification is not satisfied.

In the above passage every description is in some context; the hearer has a description of the speaker in his (the hearer's) general picture of the world. To this description of the speaker is tied the speaker's story. To the speaker's story are tied the descriptions of the individuals. What the hearer lacks are other descriptions of the same individuals which are tied directly to his general picture of the world.

Following this line of thought, the sentence "did you wind the clock" might cause a machine to create a structure which would look in part like Figure 6. The nodes clock-1 and the-clock-1 will be termed co-descriptors, because they both describe the same individual. The term co-description is chosen over co-reference or co-denotation because no implication of speaker intentions or existence of the individual described is intended.

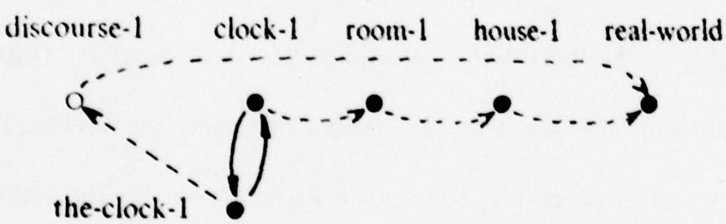

Figure 6
Co-Description Between Clock-1 and The-clock-1

In presenting Figure 6 we assume that the sentence "did you wind the clock" undergoes some sort of lexical and structural analysis. This analysis results in a description which will be termed a sentence structure. The node the-clock-1 in Figure 6 is a sentence structure co-descriptor of the words "the clock" in the sentence "did you

wind the clock". It is during the listener's process of analysis that he decides to "understand" the description the-clock-1 by making it a co-descriptor of clock-1. Had the phrase "a clock" been used, the listener might well have chosen not to find a co-descriptor of the corresponding node, a-clock-1 (particularly if he didn't wind any clocks).

In Figure 6 the relationship between clock-1 and the-clock-1 is shown to be absolutely symmetrical. Yet one intuitively feels that clock-1 is somehow more fundamental than the-clock-1. That, for example, clock-1 should be the "value" of the-clock-1. Indeed, it should be, if we view the purpose of evaluation to be the discovery or creation of more useful descriptions of entities - e.g. 2 + 2 being for most purposes a less useful description than 4. It is the description clock-1 which is already tied through sense data to the world. Clock-1 is the most useful description for most purposes (but not, for example, for understanding the anaphoric references discussed by Webber). The difference between clock-1 and the-clock-1 is captured in Figure 6 not by the links between them but by the structures in which they occur. Procedures which operate on such data structures must have the goal of finding co-referent descriptions which meet certain criteria, such as being a single number, or being a role in a certain structure. It is assumed that data structures like the one in Figure 6 will be operated on by procedures which are either implicitly or explicitly so goal directed.

3.2 Nodes in Sentence Structure

To establish the network in Figure 6 the hearer must be able to make the inference that in discourse-1, the-clock would refer to clock-1. But surely not every occurrence of the-clock in discourse-1 will refer to clock-1. Rather, it is the only occurrence in an instance of the sentence "did you wind the clock". Figure 6 needs to be amended to Figure 7. Here the-clock-1 is an instance of the-clock which occurs in an instance of did-you-wind-the-clock which occurs in discourse-1.

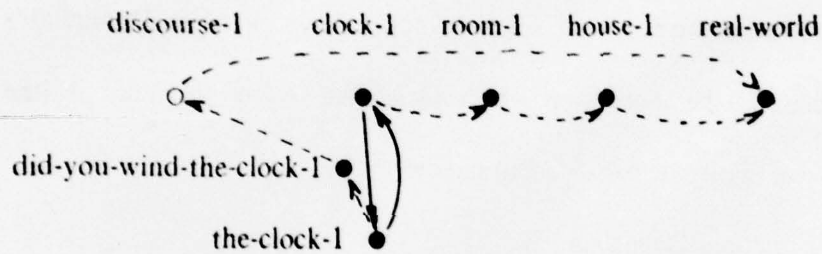

Figure 7

In her study of anaphora, Webber (1978) also points out the importance of recording the sentence in which an expression was used. Consider the sentences:

3.1) Wendy bought a yellow T-shirt that Bruce had liked.
3.2) It cost twenty dollars.

Webber points out that an appropriate description of the entity referred to by "it" in the second sentence is not "the yellow T-shirt that Bruce had liked" since sentence 3.1 is true even if Bruce had liked several T-shirts (and both the speaker and the listener were aware of the fact). Nor is it "the yellow T-shirt that Bruce had liked and Wendy

bought", since sentence 3.2 can be true even if Wendy had bought several such T-shirts. What is an appropriate description for this entity is something like "the yellow T-shirt that Bruce had liked and that Wendy bought and that was mentioned in sentence 3.1." This interpretation is achieved by a structure like that in Figure 7.

3.3 <u>Referential and Attributive Use of Definite Descriptions</u>

Donnellan (1966) has suggested that definite descriptions like "the clock" have two uses, referential and attributive. What we have done above captures Donnellan's referential use. Fortunately, the distinction which Donnellan wants to make is one which our development can easily be expanded to handle.

Explaining his distinction, Donnellan says

> I will call the two uses of definite descriptions I have in mind the attributive use and the referential use. A speaker who uses a definite description attributively in an assertion states something about whoever or whatever is the so-and-so. A speaker who uses a definite description referentially in an assertion, on the other hand, uses the description to enable his audience to pick out whom or what he is talking about and states something about that person or thing. In the first case the definite description might be said to occur essentially, for the speaker wishes to assert something about whatever or whoever fits that description; but in the referential use the definite description is merely one tool for doing a certain job -- calling attention to a person or thing -- and in general any other device' for doing the same job, another description or a name, would do as well. In the attributive use, the attribute of being the so-and-so is all important, while it is not in the referential use.
>
> To illustrate this distinction, in the case of a single sentence, consider the sentence, "Smith's murderer is insane." Suppose first that we come upon poor Smith foully murdered. From the brutal manner of the killing and the fact that Smith was the most lovable person in the world, we might exclaim,

"Smith's murderer is insane." I will assume, to make it a simpler case, that in a quite ordinary sense we do not know who murdered Smith (though this is not in the end essential to the case). This I shall say, is an attributive use of the definite description.

The contrast with such a use of the sentence is one of those situations in which we expect and intend our audience to realize whom we have in mind when we speak of Smith's murderer and, most importantly, to know that it is this person about whom we are going to say something.

For example, suppose that Jones has been charged with Smith's murder and has been placed on trial. Imagine that there is a discussion of Jones's odd behavior at his trial. We might sum up our impression of his behavior by saying, "Smith's murderer is insane." If someone asks to whom we are referring, by using this description, the answer here is "Jones." This, I shall say, is a referential use of the definite description.

That these two uses of the definite description in the same sentence are really quite different can perhaps best be brought out by considering the consequences of the assumption that Smith had no murderer (for example, he in fact committed suicide). In both situations, in using the definite description "Smith's murderer," the speaker in some sense presupposes or implies that there is a murderer. But when we hypothesize that the presupposition or implication is false, there are different for the two uses. In both cases we have used the predicate "is insane," but in the first case, if there is no murderer, there is no person of whom it could be correctly said that we attributed insanity to him. Such a person could be identified (correctly) only in case someone fitted the description used. But in the second case, where the definite description is simply a means of identifying the person we want to talk about, it is quite possible for the correct identification to be made even though no one fits the description we used. We were speaking about Jones even though he is not in fact Smith's murderer and, in the circumstances imagined, it was his behavior we were commenting upon. Jones might, for example, accuse of saying false things of him in calling him insane and it would be no defense, I should think, that our description, "the murderer of Smith," failed to fit him.

It is, moreover, perfectly possible for our audience to know to whom we refer, in the second situation, even though they do not share our presupposition. A person hearing our comment in the context imagined

might know we are talking about Jones even though he does not think Jones guilty.

To recapitulate, attributive use asserts something about whoever or whatever is the so-and-so. Referential use enables the user to pick out whom or what is being talked about. The distinction is brought out nicely by two sentences used by Moore (1973).

3.3) The President has been married since 1945. (*referential*)
3.4) The President has lived in the White House since 1800. (*attributive*)

Sentence 3.3 refers to the person who is currently President, while sentence 3.4 says that since 1800 it has been true of whoever was President that he lived in the White House. Both of these sentences have both attributive and referential readings, but in each case the listener can reject one reading as counterfactual.

3.4 Procedural Interpretation of Descriptions

In Figures 8 and 9 our representation has been expanded to capture the distinction between attributive and referential use <5>. The only structural difference between these two figures is that in the referential use, (Figure 8) the-President-1 is co-descriptive with a role, President-1, in the discourse, while in the attributive use, (Figure 9), it is not. Computationally the decision for referential over attributive interpretation amounts to the decision to find or create a role which is a co-descriptor of a role in the sentence. In order for this structural

difference to produce the difference in interpretation noted by Donnellan, it is necessary to introduce a specific procedure for determining the truth of descriptions.

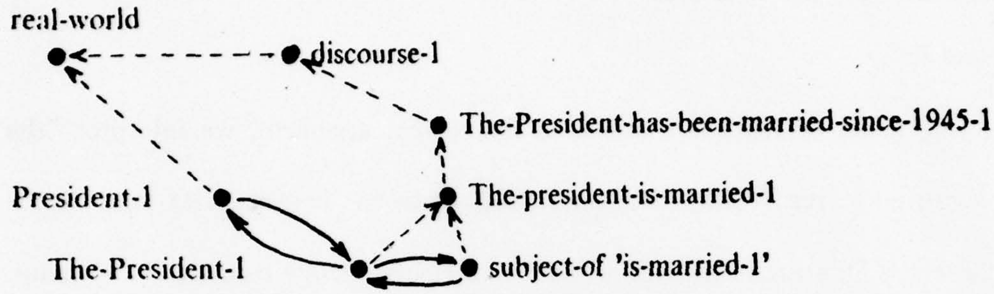

Figure 8
Referential Use of <u>the President</u>

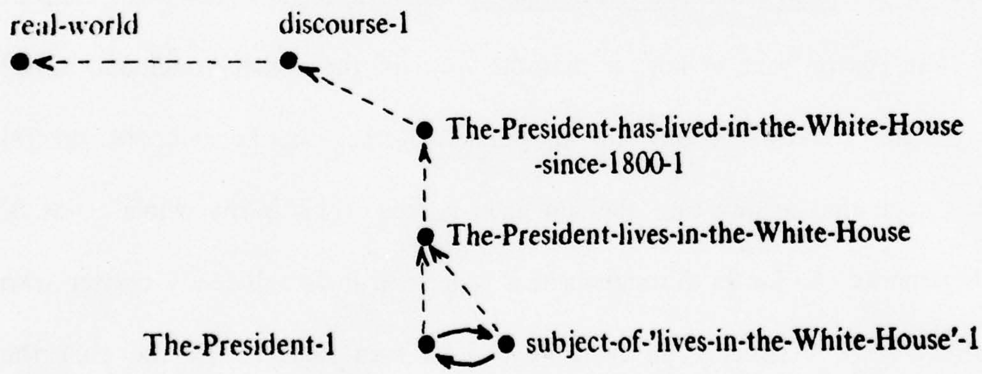

Figure 9
Attributive Use of <u>the President</u>

We will refer to an individual node as <u>true</u> if it can be put into correspondence with exactly one individual in the world; otherwise it is false. This notion of truth as correspondence with reality is an old one (see Eaton 1924) although it has long been eclipsed by the school of thought which restricts truth to propositions. We stipulate as

a part of our theory that the semantic interpretation of a description proceeds top down, as a structural decomposition. This will produce the desired interpretation of Figures 8 and 9.

Ignoring issues of tense not crucial to the current argument, we interpret "the President has lived in the White House since 1800" to mean "between 1800 and now x is true" where x is "the President lives in the White House". Now comes the key point. This is true if we can pick any time between 1800 and now and "the President lives in the White House" is true at that time. For this to be true it must be possible to pick referents for the-President-1 and the-White-House-1 which have the right relationship at that time. The central fact to note is that the truth of the higher statement is not dependent on the behavior of any one individual. Rather the behavior of several individuals is such that at any time the statement is true. This is the whole point of Donnellan's remarks. So far as this statement is concerned it doesn't really matter who is filling the role of President. In fact, we can go even farther. In the past the President has been killed and it takes a while to swear in the new Presiident. During this time there is no President, yet 3.4 is still true. It seems that 3.4 is true if there usually is a President, and when there is, he lives in the White House.

In contrast to the attributive case just covered, the referential case, "the President is married" has a role the-president-1 which must be a co-descriptor of the node president-1 which is not in its structure. The referent of president-1 is picked in

matching discourse-1, not in matching The-President-is-married-1. Therefore, we are not at liberty to re-pick it at each point in time <6>.

In summary, the semantic interpretation of a description proceeds top down, as a structural decomposition. Constraints occur when one of the roles of a description is co-descriptive with a role of some other description and thus not free to be arbitrarily chosen.

It is instructive to compare this with Moore's (1964) treatment of these two sentences using a notation based on lambda binding. Letting T-A-T name a predicate "true at time" Moore represents these sentences as

3.5) The President has been married since 1945.

```
3.6) <<LAMBDA (?X) (T-A-T (MARRIED ?X)
        <EVERY (?T) (AFTER ?T 1945)>)>
      <THE (?Y) (PRES ?Y)>>
```

3.7) The President has lived in the White House since 1800.

```
3.8) <T-A-T
        (LIVE-IN <THE (?X) (PRES ?X)> W-H)
        <EVERY (?T) (AFTER ?T 1800)>>
```

Expression 3.8 says that for every time in the set specified by <EVERY (?T) (AFTER ?T 1800)> the expression <THE (?X) (PRES ?X)> has a unique value, say X, such that (LIVE-IN X W-H) is true at that time.

Now in the case of 3.5 we want to pick the President at the current time and then use this value in an expression like 3.8. That is, in 3.8 we pick the time first, but

in 3.6 we want to pick the President first. Since this notation, like LISP, is evaluated outside in, the President is picked first by lambda abstracting it to an outer expression.

We achieve by a pointer to a co-descriptive expression what Moore achieves by lambda abstraction. Is there a significant difference? We claim that our notation should be more effective in minimizing the difference in the representation and interpretation of the two cases. In our notation one need not worry about the referential/attributive distinction until the the-President node is reached. Also in our notation the referential case is a strict addition to the attributive case. This is important because in most sentences, e.g. "The president has owned a terrier since 1977" the listener will not be able to distinguish between the two readings. If a later sentence requires the listener to make a distinction, a minimal alteration to his data structures is required. Also, the degree of semantic distinction between the two readings depends on how much is known about the discourse co-descriptor. This addresses an issue mentioned in Partee (1972)

> "having a particular individual in mind (the 'referential' case) and knowing nothing about an individual other than some descriptive phrase (the 'attributive' case) may be just two extremes on a continuum of 'vividness'. One may consider, for instance, the case of a detective tracking down a criminal and obtaining more and more clues, including fingerprints, voice recordings, photographs of varying clarity, etc. It is not at all clear at what point the detective, who may be described as 'looking for the man who did so-and-so' stops looking for 'whoever it is that did so-and-so' and starts looking for a particular individual."

3.5 Subject Co-Descriptors

Both Figures 8 and 9 employ co-description between a the-President node and a subject-of-x node in order to indicate the role played by the noun phrase in the verb. More evidence for this structure can be supplied in the form of an ambiguity brought out by verb phrase deletion.

Sag (1976) gives a number of examples to argue that verb phrase deletion occurs under identity of structure. For example, the following sentences are two, not four ways ambiguous.

3.9) John likes flying planes, and Bill does too.
3.10) The chickens are ready to eat, and the children are too.

Now consider the following sentence (Partee 1978)

3.11) The prosecutor believed that he would win the case, and so did the defense attorney.

The missing verb phrase can be understood in two ways. Either the defense attorney believed that he, the defense attorney, would win, or the defense attorney believed that the prosecutor would win. Assuming some version of deletion under identical structure to be operative, and assuming referential intepretation, these two readings argue that the phrase "the prosecutor believed that he would win" can have either of the structures shown in Figures 10 and 11. In Figure 10, he-1 is co-descriptive with the subject role, so when the defense attorney replaces the prosecutor as a co-descriptor of the subject, the co-descriptor of he-1 changes. This structure has been expressed with lambda calculus as

λ(x) (x believed that x would win).

In Figure 11, on the other hand, he-1 is co-descriptive with the-prosecutor-1, and can stay so even when the-defense-attorney-1 replaces the-prosecutor-1 as co-descriptor of the subject. Partee (1978) has termed this pragmatic anaphora, that in Figure 10, bound variable anaphora.

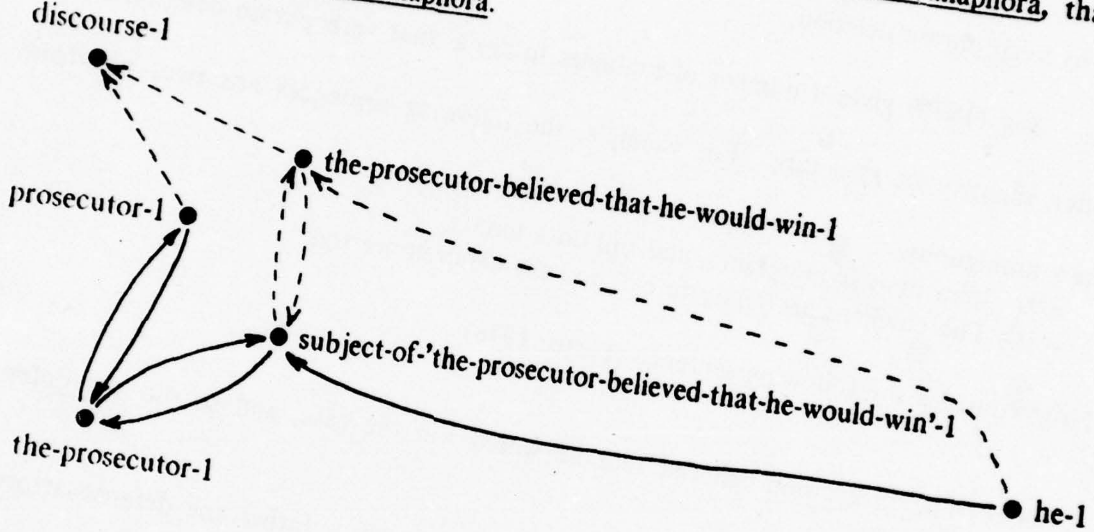

Figure 10

Structure of "the prosecutor believed that he would win" leading to "the defense attorney believed that the defense attorney would win".

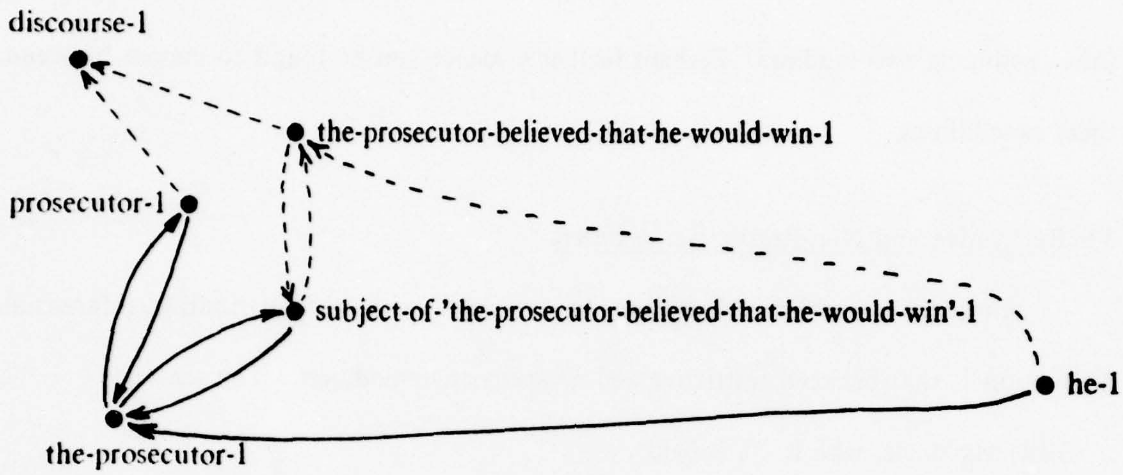

Figure 11

Structure of "the prosecutor believed that he would win" leading to "the defense attorney believed that the prosecutor would win".

It is fair to ask whether, in Figure 11, he-1 shouldn't be co-descriptive with the discourse role prosecutor-1 instead of the sentence role the-prosecutor-1. The trouble with this is that the ambiguity also occurs with attributive interpretations where there is no discourse role. Consider, for example:

3.12) Since 1800, the President has believed that he was the top government figure and so has the Chief Justice of the Supreme Court.

Another, more attractive, possibility would be to claim that

3.13) The prosecutor believed that he would win.

has the single structure shown in Figure 11, and that verb phrase deletion would optionally copy the co-description with he-1 to the new co-descriptor of the subject role,

thus producing two readings. Perhaps further evidence can be found to choose between these possibilities.

3.6 Restrictive and Non-Restrictive Modifiers

Another opposition which can be captured with the attributive/referential distinction is that between restrictive and non-restrictive modifiers. The sentence

3.14) My uncle, who is 70, is bald.

has a non-restrictive relative clause, "who is 70", modifying, "my uncle". That is "who is 70" is not used to pick out the uncle who is bald, but just gives extra information about him. This sentence is equivalent to the two sentences

3.15) My uncle is 70.
3.16) My uncle is bald.

By contrast, in the restrictive reading of

3.17) My uncle who is 70 is bald.

"who is 70" is used to pick out a particular uncle.

The distinction between 3.14 and 3.17 can be captured as shown in Figures 12 and 13.

Note that the head of the relative clause must be either attributive or referential for both main and relative clauses. It is not possible to have, for example.

3.18) The President, who has been married since 1945, has lived in the White House since 1800.

This argues that the relative pronoun, here "who", is a co-descriptor of a sentence role of the main clause, not a discourse role.

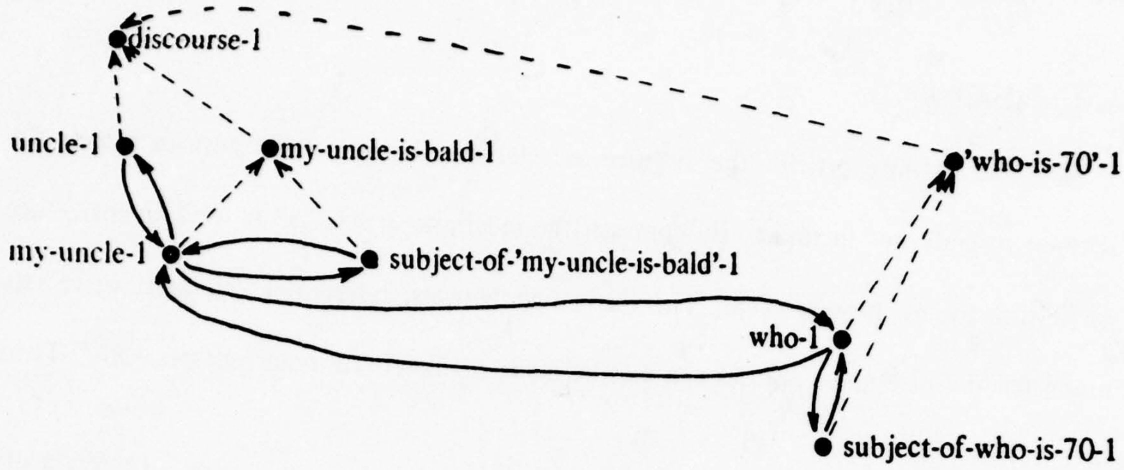

Figure 12
My uncle, who is 70, is bald.

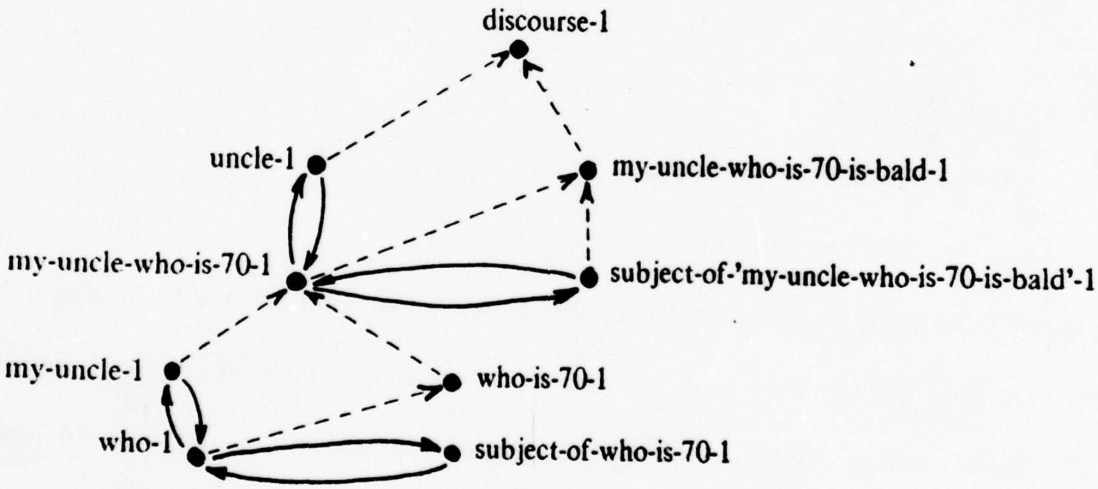

Figure 13
My uncle who is 70 is bald

4. Indefinites

The definite article "the" is used in references to boundable entities which the listener already has in mind. By contrast, the indefinite article "a" is used to introduce an entity to the listener <7>. The distinction between referential and attributive use made for definite noun phrases can also be made for indefinite noun phrases <8>. Thus in parallel with

3.3) The President has been married since 1945. *(referential)*
3.4) The President has lived in the White House since 1800. *(attributive)*

we have

4.1) A member of the House has been married since 1945. *(referential)*
4.2) A white man has occupied the White House since 1800. *(attributive)*

4.1 Opaque Operators

Quine (1960) pointed out that

4.3) Tom believes that someone is a spy.

can mean either i) that there is a particular person whom Tom believes to be a spy, or ii) that Tom believes there is at least one spy in the world. Following the conventions set forth above, these would be represented as shown in Figures 14 and 15. Figure 14 may be interpreted to mean that a particular person, P, who the speaker does not expect the listener to identify, has been referred to by "someone" and that Tom believes that someone is a spy is true when person, P, is the referent of "someone". Figure 15 says only that Tom believes that someone is a spy is true.

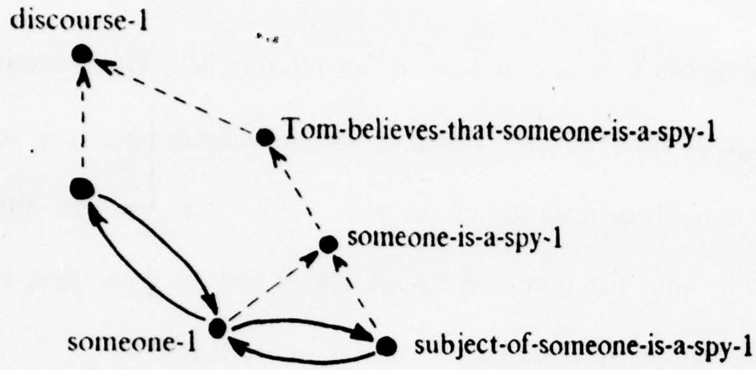

Figure 14
(where Tom knows who the spy is)

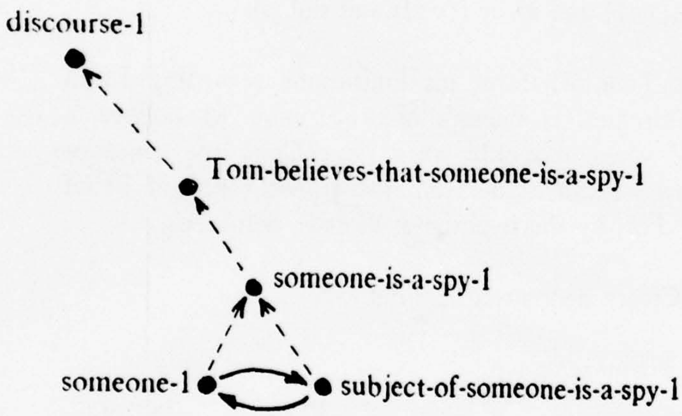

Figure 15
(where Tom believes a spy exists)

When he presented these sentences Quine refered to the reading of Figure 14 as the transparent sense of belief and the reading of Figure 15 as the opaque sense of belief. He represented these two readings using the existential quantifier as:

$(\exists x)$(Tom believes that x is a spy). *(transparent)*

Tom believes that $(\exists x)(x$ is a spy). *(opaque)*

The transparent reading was so-called because it allowed "quantifying in". Quine went on to point out that this transparent / opaque dichotomy was an unacceptable way to distinguish between the above two readings of believe. This was because the unrestricted ability to quantify into the transparent sense allows one to show that if Tom believes anything in the transparent sense then he believes everything.

Quine explained the difficulty with transparent belief as follows:

> Where '*p*' represents a sentence, let us write 'δp' (following Kronecker) as short for the description:
>
> the number x such that ((x=1) and p) or ((x=0) and not p).
>
> We may suppose that poor Tom, whatever his limitations regarding Latin literature and local philanthropies, is enough of a logician to believe a sentence of the form 'δp=1' when and only when he believes the sentence represented by '*p*'. But then we can argue from the transparency of belief that he believes everything. For, by the hypotheses already before us,
>
> (3) Tom believes that δ(Cicero denounced Catiline)=1.
>
> But, whenever '*p* represents a true sentence,
>
> δp=δ(Cicero denounced Catiline).
>
> But then, by (3) and the transparency of belief,
>
> Tom believes that δp=1,
>
> from which it follows, by the hypothesis about Tom's logical acumen, that
>
> (4) Tom believes that *p*.
>
> But '*p*' represented any true sentence. Repeating the argument using the falsehood 'Tully did not denounce Catiline' instead of the truth 'Cicero denounced Catiline', we establish (4) also where '*p* represents any falsehood. Tom ends up believing everything.

> Thus in declaring belief invariably transparent for the sake of (2)[1] and 'There is someone whom I believe to be a spy', we would let in too much. . . . In general what is wanted is not a doctrine of transparency or opacity of belief, but a way of indicating, selectively and changably, just what positions in the contained sentence are to shine through as referential on any particular occasion.

Quine went on to suggest a way of doing this which is closely related to ours. We settle on a single sense of belief, the opaque sense, and think of this as a relation between a believer and an intension named by a <u>that</u> clause. Intensions named by such clauses without free variables -- intensions of degree 0 -- are propositions. There are also intensions of degree 1, or attributes, named by prefixing a variable to a sentence in which it has free occurrence; for instance z(z is a spy) is the attribute of spyhood. We name higher degree intensions by prefixing multiple variables in the same fashion. Using this notation we can write

($\exists$x)(Tom believes y(y is a spy) of x).

Quine used this notation to single out exactly the positions which were to be referential just as we have done in Figure 14. Our solution and his are based on the same idea of singling out referential positions and will share many of the same properties. <9>

It remains to examine what happens when opaque contexts are nested. Consider the sentence.

4.4) John told me that Carol said that <u>the man who killed Kennedy</u> wore size twelve galoshes.

---

1. 'Someone is such that Tom believes that he denounced Catiline.'

As before, the phrase "the man who killed Kennedy" can be either attributive or referential. If it is referential, then "the man who killed Kennedy" may be Carol's words, or John may have been paraphrasing Carol's reference, or the speaker may have been paraphrasing John. Thus we have one attributive interpretation and three versions of the referential one - depending on whom we credit with the expression, "the man who killed Kennedy". Note that in

4.5) John told me, "Carol said that <u>the man who killed Kennedy</u> wore size twelve galoshes."

the quotation marks indicate that the speaker has not paraphrased John's statement, and so we are left with only two versions of the referential interpretation, along with the attributive one.

The most natural way for us to represent these facts is to set up explicit representations of the discourses between John and Carol and between the speaker and John. Doing this, the three interpretations of 4.5 would be as indicated in Figures 16, 17, and 18. Note that in these figures <u>discourse-1</u>, the discourse containing our sample sentence, is in the real world while <u>discourse-2</u> and <u>discourse-3</u> are in <u>discourse-1</u>. The listener is only hearing about them indirectly. Consider the case, shown in Figure 17, where "the man who killed Kennedy" is John's paraphrase of what Carol said. In this case <u>the-man-who-killed-Kennedy-1</u> is a co-descriptor of <u>man-2</u>, a man role in the discourse between John and the speaker.

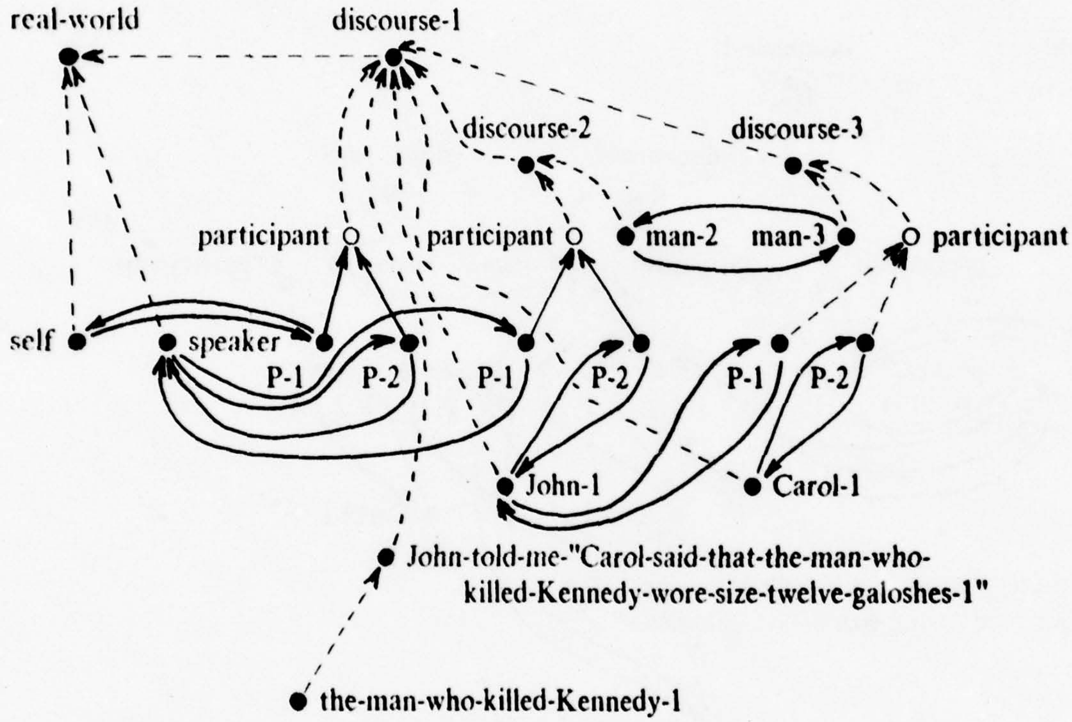

Figure 16

Attributive interpretation of "John told me 'Carol said that the man who killed Kennedy wore size twelve galoshes'."

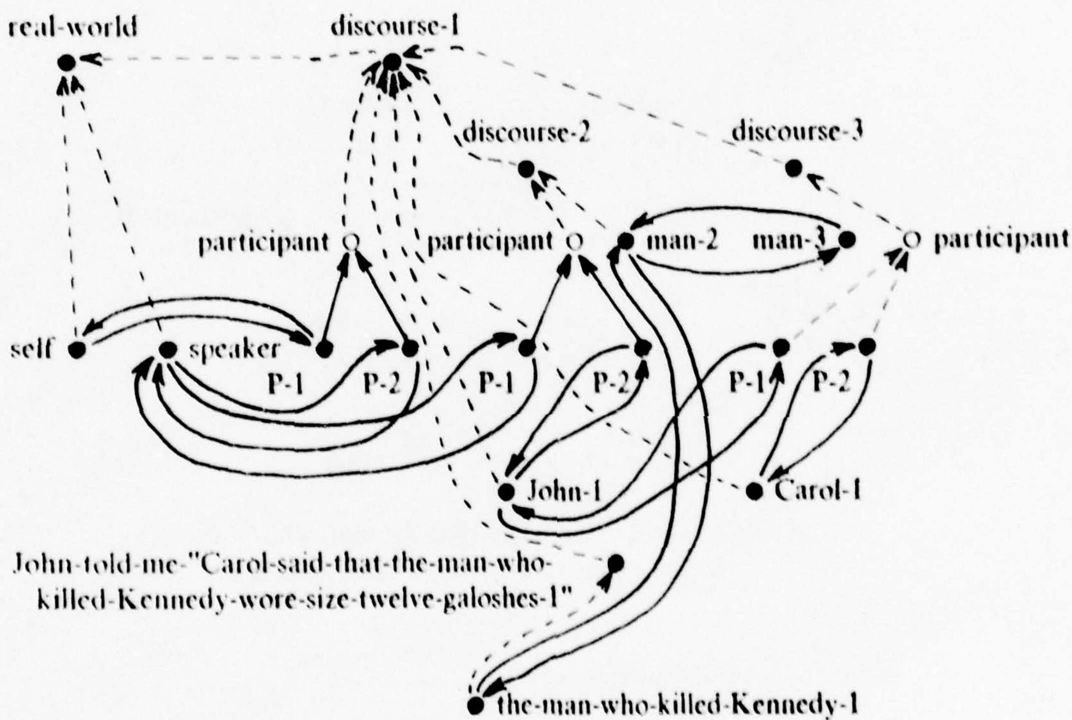

Figure 17

(Where John paraphrases Carol)

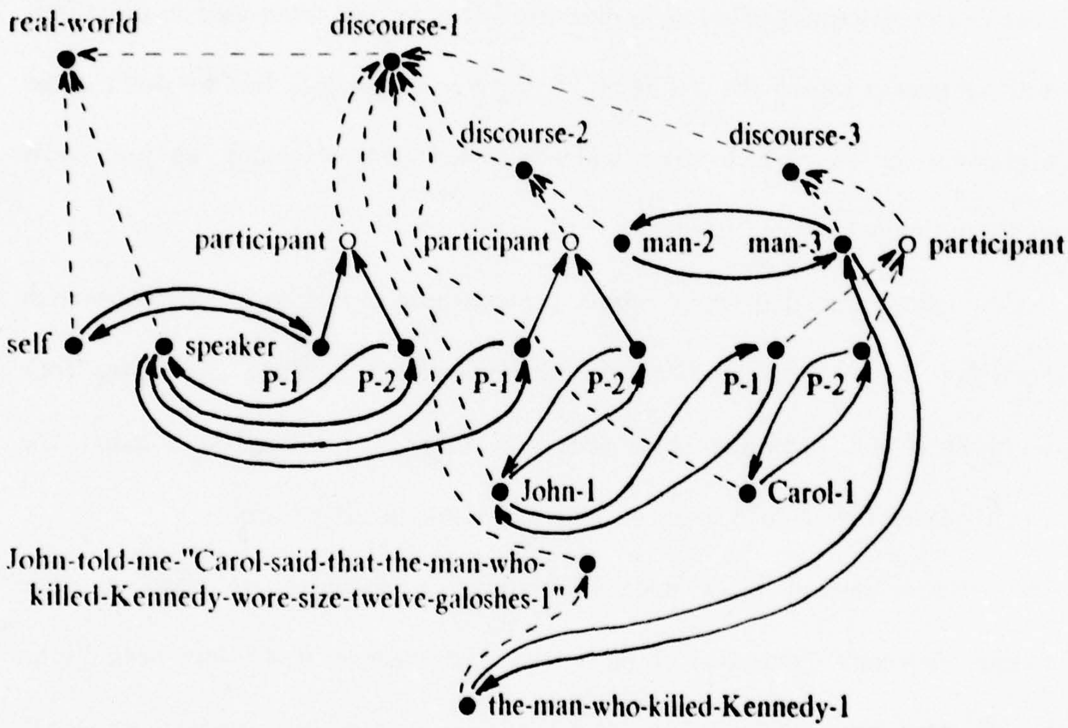

Figure 18

(where John quotes Carol)

This means that the-man-who-killed-Kennedy-1 must be matched to the world subject to the constraint that it matches the same entity as man-2, a role analogous to one which was presumably set up during the discourse between John and the speaker. If the phrase is a co-descriptor of a role in discourse-2 it must have been used in discourse-2. Carol's statement caused the setting up of man-3 in discourse-3, but we don't know what expression she used which was a sentence co-descriptor of man-3, we just know that man-2 and man-3 are co-descriptors. <10, 11>.

In summary, note that nested opaque contexts have caused us to part ways with the quantifier scope interpretation of opaque contexts. While we give four interpretations of 4.4, quantifier scope gives only three - depending on whether the quantifier is placed before John, between John and Carol, or after Carol.

Sometimes, instead of a node representing a discourse we need a node representing someone's description of the world. For example, if 4.5 had been "John told me, 'Carol believes...'" then discourse-3 would be replaced by a node representing Carol's model of the world.

5. Discourse Iteration

In the previous section we saw how nested opaque operators could lead to multiple versions of the referential interpretation. The same thing can happen when quantifiers are introduced. Consider the sentence

5.1) Every boy wants a lion.

We take this sentence to be ambiguous between

i) They will take any lion.
ii) Each wants his own specific lion.
iii) They all want the same specific lion.

These three readings can be expressed as shown in Figures 19-21.

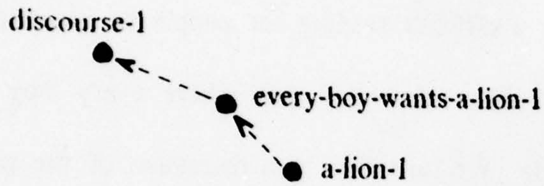

Figure 19
They will take any lion

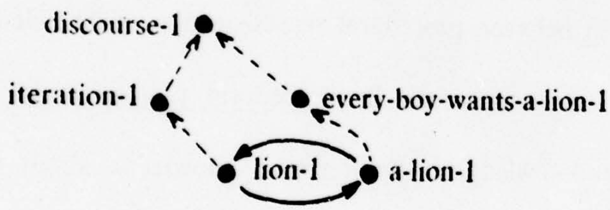

Figure 20
Each wants his own specific lion

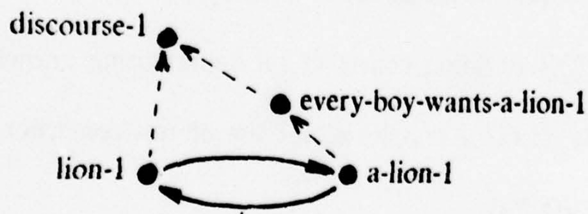

Figure 21
Each boys wants the same specific lion

Figures 19-21 can be interpreted as follows: In each case <u>want-a-lion</u> must be

true for each boy. In Figure 19, a-lion-1 is attributive. It can be picked as needed for each wanting. Figure 20 uses a new construct. Here a-lion-1 is referential, but it refers to a role, lion-1, in the iteration, iteration-1, of want at the discourse level. It's not surprising, given our account, that this is a difficult reading for people to construct. Finally, Figure 21 gives the by now familiar referential case where every boy is constrained to wanting the same lion. <10> We turn now to a discussion of the new construct used in Figure 20, discourse iteration.

5.1 Procedural Representation of Knowledge

It may be helpful to distinguish between procedural representation of knowledge and representation of knowledge about procedures. By procedural representation of knowledge we will mean casting that knowledge in the form of knowledge about the results of a procedure. For example, the knowledge that every boy wants a lion could be cast in the form "if you check every boy and count those who want lions, then the count of those who want lions will be equal to the count of the boys checked."

Woods (1977) introduced a FOR iteration construct for representing knowledge of quantified propositions procedurally <13>. Examples of the use of this construct are:

(FOR EVERY X / CLASS : (P X) ; (Q X))
  "Every X in CLASS that satisfies P also satisfies Q."
(FOR SOME X / CLASS : (P X) ; (Q X))
  "Some X in CLASS that satisfies P also satisfies Q."
(FOR GEN X / CLASS : (P X) ; (Q X))
  "A generic X in CLASS that satisfies P will also satisfy Q."
(FOR THE X / CLASS : (P X) ; (Q X))
  "The single X in CLASS that satisfies P also satisfies Q."

(FOR (ORDINAL 3) X / CLASS : (P X) ; (Q X))
"The third X in CLASS that satisfies P also satisfies Q."
(FOR (GREATER 3 E) X / CLASS : (P X) ; (Q X))
"More than 3 X's in CLASS that satisfy P also satisfy Q."
(FOR (EQUAL 3 E) X / CLASS : (P X) ; (Q X))
"At least 3 X's in CLASS that satisfy P also satisfy Q."
(NOT (FOR EQUAL 3 E) X / CLASS : (P X) ; (Q X)))
"Fewer than 3 X's in CLASS satisfy P and also satisfy Q."
(EQUAL 3 (NUMBER X / CLASS : (P X) : (Q X)))
"Exactly 3 X's in CLASS satisfy P and also satisfy Q."

Woods assumes that for each CLASS there exists an enumeration function which will produce all the members of the CLASS. The FOR statement specifies that the predicate P(x) be applied in turn to each element aduced by the enumeration function. The predicate Q(x) is then applied to those for which P(x) is true.

The first argument of the FOR divides the statements into those (EVERY, THE EQUAL GEN) which exhaust the enumeration function and those (SOME, ORDINAL, GREATER) which enumerate either until exhaustion or until a termination criterion is satisfied. The description of the termination conditions and the resulting states assumes that, in general, three counts are maintained.

i) a count, C, of the elements enumerated so far.

ii) a count, CP, of the elements passing P(x) so far.

iii) a count, CPQ, of the elements passing P(x) and Q(x) so far.

For example, the (FOR EVERY ...) statement says that in the situation resulting from its execution CP = CPQ.

The FOR iteration construct allows knowledge about sets to be built up from knowledge about individuals contained in the sets. As was mentioned in Section 2.3, such a composition operation should not be used to the exclusion of decomposition. A decompositional approach allows predicates on sets of individuals rather than just predicates on individuals as is done here; predicates on individuals are inferred from predicates on the sets they are in. Both composition and decomposition are needed.

By converting Wood's FOR construct to our notation, the difference between the two referential readings of "every boy wants a lion" can be spelled out in more detail as shown in Figures 22 and 23. Figure 22 shows the situation where each boy wants a different lion. On our account, to understand this reading it is necessary to construct a structure in the discourse in which lions are, in effect, individuated by the boys who want them. In doing this, there are two ways that "every boy wants a lion" could be understood.

i) If you pick any boy, he wants a lion. This is the GEN option of the FOR iteration.

i) If you iterate through the boys and check each, each will want a lion. This is the EVERY option of the FOR iteration.

In this example, either option could be picked. The use of the quantifier "every" (as contrasted with, for example, "any boy wants a lion") suggests, but does not force the EVERY option. <14> In Figure 22, a FOR iteration with the EVERY option has been chosen. This iteration is represented by the node FOR-EVERY-

BOY/BOYS:T;WANT-BOY-LION-1. This node has a generic role BOY which always describes the boy for the current iteration, just as the-President describes whoever is currently President. Given BOY there is always an individual wanting of a lion, WANT-BOY-LION-1, unique to BOY, and given this wanting, there is a unique lion, LION-1. This last node is made a co-descriptor of the sentence node a-lion-1.

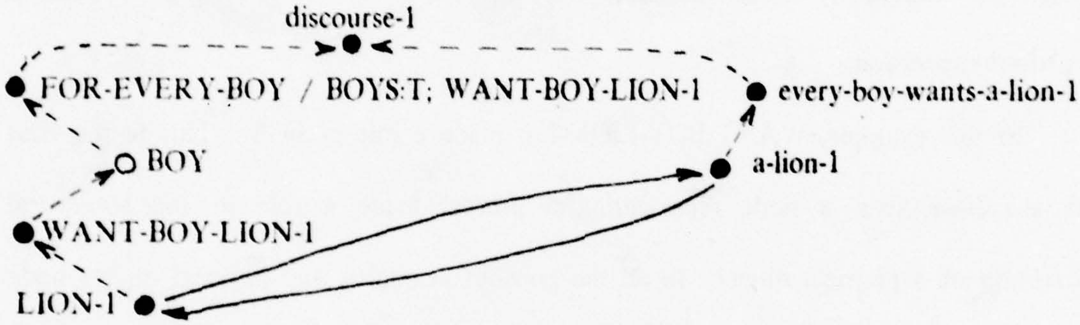

Figure 22
Each boy wants a different lion.

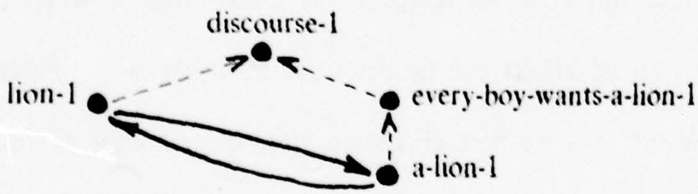

Figure 23
Each boy wants the same specific lion.

FOR-EVERY-BOY/BOYS:T;WANT-BOY-LION-1 can be matched to the world if WANT-BOY-LION-1 can be matched to the world uniquely for every possible match of BOY to the world.

The reader might note here how we have utilized Fahlman's notion of an individual only being individual to some structure (refer to Section 2.2). LION-1 is unique only within BOY. As BOY describes different boys, the properties of LION-1 are inherited to different lions. <15>. Note also that this formulation allows the same lion to be wanted by more than one boy, the standard understanding of such a quantified expression.

In this example, WANT-BOY-LION-1 is made a role in BOY. This is the first time we have seen a node representing a process made a role in the structural description of a physical object. In all the previous examples, any physical object node was made a role in a state or process node. The following question is raised:

i) If someone wants something, is that wanting an attribute of the wanter, or does the wanter fill a role in the wanting process.

Addressing a more general form of this question, Karp (1975) claims that it is all a matter of viewpoint. The same state of affairs can be described in either way. Since the lion is to be individuated by the boy, we here choose to view the wanting of the lion as an attribute of the boy.

In contrast to Figure 22, Figure 23 contains no iteration at all. Figure 23 represents the case where all the boys want the same lion. Since in this case nothing is individuated to individual boys, there is nothing to be gained by expanding the sentence structure into an iteration in discourse structure. On this reading, "every boy wants a

lion" can be conceptualized just like "a boy wants a lion", provided we treat the <u>quantity</u> of an entity (and other properties conveyed by quantifiers and determiners) as a property no different from other properties of the entity. A more standard version of this idea is that all predicates will be predicates on sets, with predicates on individuals being the special case of predicates on sets of one element. "A boy" describes a one element set, "every boy" describes a many element set. The predicate "wants a lion" can be applied to either of these, since the number of elements is not particularly important.

It is important to distinguish between

a) a predicate *on* a set which describes the set as a single entity. For example, it gives the *number of elements*.

b) a predicate on a set which describes properties of the individual elements of the set.

c) a predicate on an element of a set.

The choice between b) and c) as a method of describing individuals is made in favor of b) when we wish to combine the functions of a) and b) into a single predicate. This, I will argue in Section 6, is what often happens in English.

When a set is predicated as a single entity, I will refer to it as a <u>collective</u>. When it is predicated as both a set and individuals, I will refer to it as a <u>plurality</u>. The reader may notice that I am at pains here to undo the decomposition of everything into predicates on sets and individuals long favored by mathematicians. I do this, as I argue in Section 6, to avoid a forced decomposition of predicates on pluralities.

5.2 The Collective / Distributive Distinction

In applying predicates to pluralities we must be careful not to confuse a predicate on a plurality with a predicate on a collective. If a process, like wanting, is done by a plurality then a plurality of wantings occurs. This contrasts with a collective, which acts as a single unit. Fauconnier (1975) demonstrates this distinction with the examples

The men gathered / united / quarreled.	collective
The men took off their hats.	distributive
The men carried the couch.	ambiguous

As a further example, observe that in

5.2) Everybody gave \$1000 to many of the men.

we must decide whether the men receiving money receive \$1000 apiece or just participate in a group receiving \$1000. We must also decide whether "everybody" acted individually or collectively.

Note that one cannot personally receive \$1000 if \$1000 is given to a group he is in, but one can be personally told the news if the news is told to a group he is in. From this we see that the distributive reading may or may not require a separate instance of some state or process for each individual.

A plural noun group may be understood to be either collective or distributive. If distributive, each individual may be treated on an individual basis, or as part of a group, e.g. "I told all the boys the news" may mean one telling (collective), one telling per boy,

or some number of tellings less than one per boy. Vendler (1967) gives examples intended to show that the choice between "each" and "every" influences our preference for the number of occurrances.

> "Suppose I show you a basket of apples and I tell you
>
> Take all of them.
>
> If you started to pick them one by one I should be surprised. My offer was sweeping: you should take the apples, if possible, 'en bloc'. Had I said
>
> Take every one of them.
>
> I should not care how you took them, provided you do not leave any behind. If I say
>
> Take each of them.
>
> one feels that the sentence is unfinished. Something like
>
> Take each of them and examine them in turn.
>
> is expected. Thus I expect you to take them one after the other not missing any."

"All" favors collective interpretation but takes distributive. "Every" favors distributive interpretation. "Each" strongly favors distributive interpretation and favors distinctive actions for each individual.

5.3 Syntactic Structures and Iteration Loops

Suppose there is a noun phrase NP_i to be iterated over and a second noun phrase $NP_{a/r}$ which can be either attributive or referential. We can identify three ways these can be related syntactically.

a) NP_i contains $NP_{a/r}$ as a syntactic substructure.

every flight to an eastern city
every flight which was going to an eastern city.

b) $NP_{a/r}$ contains NP_i as a syntactic substructure.

a guy who is representing each raw rubber producer in Brazil
a guy representing each raw rubber producer in Brazil
a representative of each raw rubber producer
each raw rubber producer's representative
a representative of the manager of the critical section of each stream

c) Neither is a syntactic substructure of the other

i) Each secretary reminded me about the scheduling of an appointment.
" about scheduling an appointment.
" to schedule an appointment.
" that I should schedule an appointment.
Mary intends to mail each of her suicide notes to a friend.
" a friend each of her suicide notes.

ii) Every boy wants a lion.
Each teacher thinks that a student of mine was called before the dean.
Each teacher overheard a rumor that a student of mine was called before the dean.
Each teacher overheard the rumor that a student of mine has been called before the dean. <16>

Case a. is handled by deciding whether $NP_{a/r}$ is attributive or referential. If attributive it is picked as needed, if referential there is one entity referred to. Nothing is gained by constructing an iteration loop.

Cases b. and c. are different. Whether $NP_{a/r}$ is picked as attributive or referential the question remains whether NP_i is taken as collective or distributive. For

example, if each raw rubber producer in Brazil is taken as collective, then a single guy represents this collective entity. If it is taken as distributive, then there is more than one representing and thus possibly more than one guy. The four possibilities are:

a) attributive, collective
There is one guy.

b) attributive, distributive
There is more than one guy.

c) referential collective
There is one specific guy, he has a co-descriptor in discourse structure.

d) referential distributive
There can be either one specific guy doing multiple representings, he has a co-descriptor in the discourse structure; or there can be a specific guy for each producer, he has a co-descriptor in an iteration loop.

The decision between a collective and a distributive reading is influenced by the syntactic structure, the semantics of the lexical items involved, and the varying distributive force of words like "each", "every", and "all". Van Lehn (1978) reports the results of experiments which show wide variations in judgement among subjects.

In a phrase like

5.3) a representative of the manager of the critical section of each stream

one first asks whether "the critical section of each stream" is collective or distributive. That is, whether or not there is one critical section for all the streams. Assuming several critical sections, one then asks whether there is one manager. If one assumes several managers, the question of how many representatives comes up. This shows how

the collective/distributive distinction must be made working outward from the inner most level of nesting containing a potentially distributive quantifier. At each level both syntactic and semantic evidence is used to make the decision. <17> Van Lehn (1978) described a procedure similar to this - he called it quantifier raising - which he felt gave the best approximation to his data on human processing of quantifiers, of the quantifier interpretation methods known to him. Fodor (1979b) also proposed a method like this to account for human processing of quantifiers.

5.4 Multiple Iterations

The previous sections suggested that an iteration need be constructed only when two descriptions in the sentence were related in a special way, e.g. one individuated by the other. Otherwise, predicates on "pluralities" provide an adequate description. For example,

All the boys kissed all the girls.

could be represented by a double iteration, but in general we claim this is not necessary. More arguments in support of this claim will be advanced in Section 6. One case where construction of an iteration is useful is when an apparent double iteration is actually best accomplished by a single iteration! Consider the sentences

5.4) Each cork is fastened to each bottle by a small wire basket.

5.5) Each man and each woman will be joined in marriage here tonight.

Most people prefer to understand sentence 5.4 by iterating over bottles and individuating the corks by the bottles. In sentence 5.5 they iterate over man/woman pairs.

Note that these sentences provide evidence that a quantifier like "each" doesn't necessarily set up an iteration. It usually does, and therefore these sentences are a bit odd. But these sentences are understood when each distributed quantified expression has a discourse co-description which participates in an iteration.

Setting up an iteration in 5.4 and 5.5 expresses the specific dependence of corks on bottles, and men and women on each other.

5.5 Constructing an Iteration

In the previous example we could have had each man married to each woman! Indeed, it is not always obvious how an iteration should be constructed. When someone hears

> 5.6) A requirement for the course is the carving of a block of wood into each of the 12 designs.

his reasoning might be as follows: "Well, let's see. We take the wood and carve the first design. (Pursuing the distributive referential reading with one block) Oh! Oh! Now how do we carve a second design, the block is used up. Well, maybe we could fit the twelve designs on one block, or we could cut the block into twelve pieces. Or maybe I should abandon the referential reading and use twelve blocks." Here there is a real issue of how this carving can be done.

There can be no doubt that world knowledge is required to choose between readings. Consider

5.7) Everybody at MIT knows a dialect of LISP.

5.8) Everybody at IJCAI knows a dialect of LISP.

Everyone at a university knows the same specific dialect while everyone at a world conference knows different specific dialects.

In summary, we have shown how the possible readings requiring iteratio . an be represented, but not how to choose between them.

6. Pluralities

The sentence

6.1) Three lookouts saw two boats.

can be expressed in Wood's (1977) notation as

```
(AND (EQUAL 3 (NUMBER LOOKOUT/LOOKOUTS: T ;
               (FOR SOME BOAT/BOATS: T ; (SEE LOOKOUT BOAT))))
  (EQUAL 2 (NUMBER BOAT/BOATS: T ;
               (FOR SOME LOOKOUT/LOOKOUTS: T ;
               (SEE LOOKOUT BOAT)))))
```

This is a reasonably complex expression, and one can question whether it is the most perspicuous form to use, especially if it is to be processed intensionally rather than evaluated extensionally. Such a form is forced if one insists on predicates on individuals. If one allows predicates on pluralities which imply multiple individual

actions, then the above expansion of the quantifiers to positions external to a predicate is not required.

6.1 Levels of Representation

The question of whether to allow predicates on pluralities can be viewed as a choice of level of representation.

As Bierwisch (1971) notes, a sentence like

6.2) The boys hit the girls.

represents a statement whose truth would be supported by hit conditions existing between individual boys and individual girls. Indeed (Nunberg and Pan, 1975)

Have the boys ever hit the girls?

requires a reporting of even a single instance. However, since we have no information about what boy/girl pairs are involved, there seems to be no sensible way to interpret this statement in terms of conditions on individuals without making arbitrary choices. Were all the boys and girls involved? Was any girl hit by more than one boy? We don't know. So while "the boys hit the girls" implies hit relationships between individual boys and girls, there is nothing to be gained by expressing it in these terms since none of the particulars about individual hit relationships are known.

Now the important thing about "the boys hit the girls" may be the relationship between the groups, not the singling out of the individuals actually involved. Bierwisch supports this notion with sentences like

6.3) The whites oppress the negroes.

6.4) The Romans destroyed Carthage.
6.5) The Chinese of the seventh century knew porcelain.

In all these sentences we have actions or attitudes by all the members of the group which may support those actually involved. Care must be taken or this group sense may be lost in a representation focussed on individuals.

For example, a typical method of representing "the boys hit the girls" in terms of individual hittings would be to form a predicate calculus expression reading:

i) let B be a subset of the boys
ii) let G be a subset of the girls
iii) then for all members b of B there exists a girl g in G such that b hit g.
iv) and for all members g of G there exists a boy b in B such that b hit g.

In this notation the importance of the larger sets is lost. The notion that the larger set the boys was responsible for or supportive of the action of the subset, B, is gone. This could be stated separately or recovered by inference, but we take the position that it is a counter-productive to disaggregate the information in the first place. No double statement or inference should be required.

Recall the sentence "the President has lived in the White House since 1800" and how this is verified independently of what individuals are involved. We don't care who the individuals are, but only that individuals act so that such a thing is true. The same can be said for the sentences above involving plurals.

7. Reference to Generics

Sometimes "boys love candy" is represented as (∀ boy)(love (boy,candy)). This is not right. "Boys love candy" is true even if there is some particular boy who doesn't like candy. "Boys love candy" means that the typical boy loves candy, not "all boys love candy." Since we have already introduced the notion of a generic node, the proper interpretation of this sentence is obtained in our notation by making boys co-referential to the generic boy as shown in Figure 24. <18>

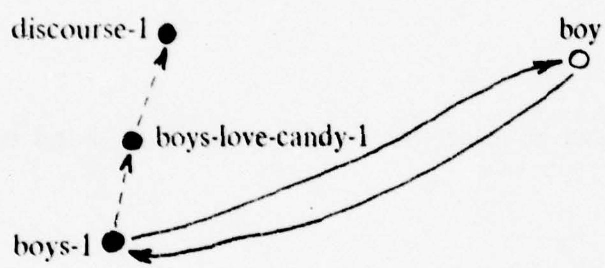

Figure 24

At least when dealing with countable entities, it is important to distinguish two alternative conceptions of universals - the conception of universals as exemplars (or generic nodes) and the conception of universals as kinds (Karp 1975). The notion that universals are exemplars might be stated as a principle of the following sort:

> (The Exemplar Principle) All and only those things are true of an abstract substance which are true of its typical examples.

So stated the principle suffers from oversimplification. As Wittgenstein (1953) points out, exemplars frequently have almost nothing common to all of them, e.g., some cups

have handles, some don't, some are glass, some are porcelain. Nevertheless, we may leave these difficulties aside for the moment and consider the simplified Exemplar Principle in distinguishing between exemplars and kinds.

The indefinite article, "a", may be used in referring to exemplars - "a lion has four legs". Exemplars may also be described using the definite article, "the", and the bare plural, e.g., "lions".

<u>A lion</u> has four legs.
<u>The lion</u> has four legs.
<u>Lions</u> have four legs.

There are other statements which can be made with "the" and the bare plural which cannot be made with "a".

<u>The lion</u> is almost extinct.
<u>Lions</u> are almost extinct.
*<u>A lion</u> is almost extinct.

<u>The lion</u> is widespread in Africa.
<u>Lions</u> are widespread in Africa.
*<u>A lion</u> is widespread in Africa.

<u>The lion</u> is a species of mammal.
<u>Lions</u> are a species of mammal.
*<u>A lion</u> is a species of mammal.

These later statements are widely recognized as referring to universals, but they ascribe to these universals properties which cannot be held by any individual lion and thus not by the exemplar. No individual lion is almost extinct, or widespread, or a species. These latter statements refer to lion kind, that is, to lions as a species. The exemplar principle is false.

Note that the sentence (Carlson 1977)

Linguists have over 30,000 books in print.

is ambiguous because "linguists" could refer either to the exemplar (or generic) linguist or to linguists as a kind. This is disambiguated by our knowledge of the world - by our knowledge that no individual linguist could write 30,000 books. On the other hand

Linguists have 62,344 legs.

can have only the exemplar interpretation, even though it might be true on a collective basis. Apparently "having legs" is something that only individuals can do while "being widespread" is something that only kinds can do.

8. Processing Considerations

In the previous sections we have introduced a new formalism for the representation of the logical form of English sentences. The new formalism is more complicated than the old. Sentences which were previously represented by predicates on individuals in combination with universal and existential quantifiers, now require a variety of mechanisms; the referential/attributive distinction, iteration loops, the collective/distributive distinction, and predicates on individuals, sets, pluralities, generics, and kinds. Two levels of representation (discourse and sentence) are used instead of one.

Why is this more complex scheme a better scheme? It can be justified in part

by the difficulties encountered in applying the quantifier scheme to certain difficult cases. For example,

i) There is no straight forward way to capture all the correct readings of nested opaque operators using quantifier scope (see Section 4.1).

ii) A sentence like

8.1) I told three of the stories to many of the men.

has three readings corresponding to

a) three stories to many men collectively
b) three particular stories to the many men distributively
c) three different stories depending on the men.

The a. reading is hard to capture unless one admits the collective/distributive distinction. <15>

iii) The "each cork fastened to each bottle" sentence requires some sort of special treatment of <u>each</u> (see Section 5.4).

These are important objections, but the primary motivation for the scheme presented here was ease of processing. There are two primary factors.

The first factor is that a representation with more constructs often provides a richer structure which can be processed more efficiently than a representation based on a smaller set of more primitive operators. I do not claim to have shown that my representation is in fact processed more efficiently, but I expect this observation to strike home with many readers familiar with the difficulties encountered by predicate calculus based theorem provers.

The second factor has to do with the desirability of delaying decisions for which the information is not yet in hand.

8.1 Human Processing of Quantified Expressions

In a recent thesis Van Lehn (1978) reports that when people are given a sentence like

8.2) A quick test confirmed that every drug was psychoactive.

they claim they understand it, but are then unable to state whether there was one test per drug or only one test for all. This ambiguity can be expressed in predicate calculus using the universal and existential quantifiers. Using the standard rule that an existentially quantified variable may only depend on the universally quantified variables to its left, we may write the two interpretations:

$(\exists \text{ test}) (\forall \text{ drug}) (C (\text{test}, \text{drug}))$
$(\forall \text{ drug}) (\exists \text{ test}) (C (\text{test}, \text{drug}))$

The difficulty with this notation is that it lacks a form which leaves the ambiguity unresolved. If people were forced to resolve the ambiguity in order to come to some understanding of the sentence they would not report to Van Lehn that they have no idea of whether there was more than one test even though they understand the sentence.

8.2 Ambiguity and Generality

The issue arises in part because of the difficulty of distinguishing between a general expression and an ambiguous one. For example, Figure 25 illustrates that both a sail boat and a motor boat can be described by the general expression boat.

This contrasts to Figure 26, where there is no concept general to the pronoun-mine and mineral-mine senses of "mine". Figure 27 points out that it is uncertain whether "a quick test confirmed that every drug was psychoactive" has a general sense. The general sense of "boat" is resolved semantically - by choosing between more specific concepts such as "sailboat" or "motorboat". If Figure 27 does have a general sense it might be resolved semantically, but it might also be resolved pragmatically - by binding the concept into the discourse structure in different ways. Philosophers have largely ignored this issue since they have been interested in formal languages whose terms are logically unambiguous. Van Lehn's results would tend to indicate that people do have a general sense of this expression.

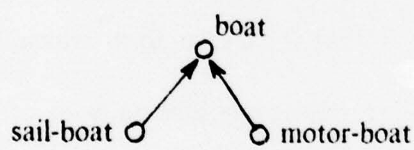

Figure 25

Figure 26

? A-quick-test-confirmed-that-every-drug was-psychoactive

$(\exists\ test)(\forall\ test)(C(test,drug))$ ○ ○ $(\forall\ drug)(\exists\ test)(C(test,drug))$

Figure 27

8.3 Processing Strategies for the Refinement of General Expressions

When a listener has both general and specific senses of a term like "boat", his choice of how specific a sense to use can be considered a processing strategy. When he hears "boat", should he choose between motor and sail boat? When he hears the "quick test" sentence, should he choose between the one and multiple test readings?

Foucounier (1975) discusses one principle which a listener might use.

> Ambiguity Principle If some particular understanding of a sentence (or consequence of it) can be treated as a presupposition in the rest of the discourse, then that understanding corresponds to a genuine reading of the sentence.

Under this principle

8.3) Frank has a boat; the sail is red with white dots.

would show that "Frank has a boat" has a reading where "boat" is a "sailboat". Further

8.4) All my friends went to see a movie last night. The one Sue and Phil saw was terrific, but the one John and Mary saw was lousy.

would imply that the first sentence had a reading where the friends saw the movie two by two.

These examples make a pretty strong case that it would be unworkable for a listener to consider all the senses of an expression implied by the ambiguity principle. This would force too much of what is called forward inferencing - the consideration of all possible interpretations of a sentence in the hopes that one will help in interpreting the following sentence, before the following sentence is available to direct the search for appropriate interpretations.

8.4 Advantages of the Representation

The advantage of the representation proposed here is that it allows the resolution of ambiguity by incrementally adding to what is already present. This is done

a) By using the attributive/referential distinction where others have used quantifier scope and making the attributive/referential distinction a question of pragmatics.

b) By making the generic/non-generic distinction a question of reference.

c) By using features to mark the distributive/collective distinction when it is desirable.

d) By using predicates on pluralities and constructing iteration loops only when needed.

All of these mechanisms work to allow easy partial interpretation of expressions which have not been fully resolved.

9. Problems and Suggestions for Further Work

The line of development proposed here has been carried far enough to allow a preliminary evaluation, but one needs to look more at, for example, verb phrase deletion, anaphora, and boolean operators.

9.1 Further Remarks on the Structural Requirements for Verb Phrase Deletion

In the sentence

9.1) Someone hit everyone.

"someone" can be understood either referentially or attributively. When it is understood referentially there is one person who does the hitting. When it is understood attributively, a different person can fill the role for each hitting.

Now a sentence like "Bill hit everyone" can only have Bill understood referentially, so it is interesting to see what happens when these two sentences are conjoined and there is verb phrase deletion. In fact, in

9.2) Someone hit everyone and then Bill did too.

someone can only be understood referentially. This argues that the structural identity required for verb phrase deletion includes the requirement that identical noun phrases be either both attributive or both referential.

Sag (1976) has presented a number of examples which might be handled with this argument. For example, there are two readings of

9.3) Betsy wants Peter to read everything Alan wants him to read.

depending on whether "everything Alan wants him to read" is taken as attributive or referential. If it is taken as attributive, then anything which happens to fit this description should be read. That is, read what Alan says. On the referential interpretation what Betsy wants Peter to read just happens to be describable as "everything Alan wants him to read". One needs to propose the details of structure and verb phrase deletion so that

9.4) Betsy wants Peter to read everything Alan does.

will have only the referential interpretation. Another similar example is

9.5) Betsy's father told her to work harder than her boss did (told her to work).

9.2 Anaphora

One might think that definite anaphoric references are made to nodes already existing in sentence or discourse structure, but the situation is more complicated. In the sentence

9.6) John got married last week and she's a great catch.

there is no explicit mention of the antecedent of she and it isn't clear when this is created or what the rules are which permit this. Note that

9.7) John is an orphan and he misses them.

doesn't work so well. The importance of this to the current discussion comes out when one considers sentences like

9.8) John traps beavers even though they are protected by law.

9.9) Senator Green believed that he had nominated the winner of the election, but Senator White believed that she had nominated him.

9.10) The alligator's tail fell off but it grew back.

In each of these sentences problems are created if one assumes that the underlined phrase and the pronoun are co-descriptive. In the first, "beavers" is non-generic and "they" is generic. In the third, a tail role and a tail are referred to. However, it isn't clear whether the phrases are co-descriptors or whether the first just makes a node which is a co-descriptor of the second accessable. <20>

9.3 Boolean Operators

Examples of the facts needing an explanation here are:

a) Conjunction of descriptions and conjunction of referents

The miner and sapper went to work.

can refer either to one person who is both a miner and a sapper or to two.

b) Conjunction and Opaque Operators

John believes that either dogs bark or cats meow.

doesn't mean the same thing as

Either John believes that dogs bark or John believes that cats meow.

c) Discourse or semantics

If it snows tomorrow, there will be chains in the truck.

could mean that chains being in the truck is conditional on snow or it could mean "this is information you will need if it snows."

d) Interaction of negation with quantification

John makes not 500 a year.

uses "not" as an adverb and means he makes less.

Every boy will not be allowed to leave.

means either that some or none will be allowed to leave, depending on whether "not" applies to the sentence or to the predicate.

10. List of Cited Works

1. Benacerraf, P. (1965) "What Numbers Could Not Be", The Philosophical Review, p.47.

2. Bierwisch, M. (1971) "*On Classifying Semantic Features*" in D. D. Steinberg and L. A. Jackobovits (eds.) Semantics, Cambridge Univ. Press, London.

3. Brachman, R.J. (1978) "On the Epistomenological Status of Semantic Networks", in N.V. Findler (ed) Associative Networks - The Representation and Use of Knowledge in Computers, Academic Press, New York.

4. Carlson, G.N. (1977) "Reference to Kinds in English" University of Mass. PhD. thesis, available from Indiana Univ. Linguistics Club.

5. Carnap, R. (1950) "Empiricism, Semantics, and Ontology" Revue Internationale de Philosophie 11 reprinted in L. Linsky (ed) Semantics and the Philosophy of Language, Univ. of Illinois Press, Urbana, Ill.

6. Cole, P. (1978) "On the Origins of Referential Opacity" in Syntax and Semantics: Pragmatics, Vol. 9 P. Cole (ed.) Academic Press, New York.

7. Donnellan, K. (1966) "Reference and Definite Descriptions", The Philosophical Review 75, 281-304.

8. Eaton, R.M. (1924) Symbolism and Truth, An Introduction to the Theory of Knowledge, Harvard Univ. Press, Cambridge.

9. Fahlman, S.E. (1979) "A System for Representing and Using Real World Knowledge", MIT Computer Science PhD thesis, to be published by MIT Press, Cambridge, Mass.

10. Fauconnier, G. (1975) "Do Quantifiers Branch" Linguistic Inquiry 6.4.

11. Fauconnier, G. (1978) "Is There a Linguistic Level of Logical Representation?" Theoretical Linguistics, Vol. 5, No. 1, Walter De Gruyter, Berlin.

12. Fillmore, C.J. (1967) "On the Syntax of Preverbs" Glossa 1, 91-125.

13. Foder, J. (1979a) submitted for publication.

14. Foder, J. (1979b) submitted for publication.

15. Goldstein, I.P. and Roberts, B.R. (1977) "NUDGE - A Knowledge Based Scheduling Program" Proceedings of 5th IJCAI, available from Dept. of Computer Science, Carnegie-Mellon Univ., Pittsburg, Pa.

16. Hayes, P.J. (1979) "On Semantic Nets, Frames and Associations", Proceedings of 5th IJCAI, available from Dept. of Computer Science, Carnegie-Mellon Univ., Pittsburgh, PA.

17. Hendrix, G. (1978) "Encoding Knowledge in Partitioned Networks", N.V. Findler (ed) Associative Networks - The Representation and Use of Knowledge in Computers, Academic Press, New York.

18. Hintikka, K.J.J. (1976) "Quantifiers in Logic and Quantifiers in Natural Languages", in S. Korner (ed) Philosophy of Logic Oxford: Basil Blackwell, 208-232.

19. Jackendoff (1972) Semantic Interpretation in Generative Grammar, MIT Press, Cambridge, Mass.

20. Karp, D.J. (1975) "General Ontology" MIT Philosophy PhD. Thesis.

21. Linsky, L (1971) Reference and Modality (ed) Oxford Univ. Press, London.

22. McCarthy, J. and Hayes, P.J. (1969) "Some Philosophical Problems from the Standpoint of Artificial Intelligence" in B. Meltzer (ed.) Machine Intelligence 4 American Elsevier Publishing Co., New York.

23. Martin, W.A. (1979) "Descriptions and the Specialization of Concepts" in P. Winston (ed) Artificial Intelligence, An MIT Perspective, MIT Press, Cambridge, Mass.

24. Minsky, M. (1975) "A Framework for Representing Knowledge" in The Psychology of Computer Vision, P.H. Winston, ed. McGraw-Hill, New York.

25. Moore, R.C. (1973) "D-SCRIPT: A Computational Theory of Descriptions" Advance Papers of the Third International Joint Conference on Artificial Intelligence, 223-229.

26. Moore, R.C. (1979) "Reasoning About Knowledge and Action" MIT Computer Science PhD thesis.

27. Nunberg, G., and Pan, C. (1975) "Inferring Quantification in Generic Sentences" Papers from the Eleventh Regional Meeting of the Chicago Linguistic Society, Chicago Linguistic Society.

28. Partee, B.H. (1972) "Opacity, Co-Reference, and Pronouns" in Semantics of Natural Language, Davidson and Harmon (eds) D. Reidel Pub. Co. Dordrecht-Holland..

29. Partee, B.H. (1978) "Bound Variables and Other Anaphors" in Proceedings of Theoretical Issues in Natural Language Proceesing - 2, available from ACM.

30. Quine, W.V.O. (1960) Word and Object, MIT Press, Cambridge, Mass.

31. Quine, W.V.O. (1970) Philosophy of Logic, Prentice-Hall, Englewood Cliffs, New Jersey.

32. Russell, B. (1905) "On Denoting" in H. Feigl and W. Sellars (eds) Readings in Philosophical Analysis, 85-102, Appleton-Century-Crofts, Inc., New York, 1949.

33. Sag, I. (1976) "Deletion and Logical Form" MIT Linguistics PhD Thesis, available through Indiana University Linguistics Club.

34. Sidner, C. (1979) "Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse", MIT Computer Science PhD Thesis.

35. Strawson, P.F. (1959) Individuals: An Essay in Descriptive Metaphysics, Anchor Books Edition, (1963).

36. Tarsky, A. (1944) "The Semantic Conception of Truth", Philosophical and Phenomenological Research, 4, reprinted in L. Linsky (ed), Semantics and the Philosophy of Language, Univ. of Illinois Press, Urbana, Ill.

37. VanLehn, K.A. (1978) "Determining the Scope of English Quantifiers" MIT Artificial Intelligence Laboratory Report AI-TR-483.

38. Vendler, Z. (1967) Linguistics in Philosophy Cornell University Press, Ithaca, New York.

39. Webber, B.L. (1978) "A Formal Approach to Discourse Anaphora", Bolt Beranek and Newman, Inc. Research Report No. 3761.

40. Wittgenstein, L. (1953) Philosophical Investigations, MacMillan, New York.

41. Woods, W. (1977) "Semantics and Quantification in Natural Language Question Answering" Bolt Beranek and Newman Report 3687.

Notes

1. The current thinking of SRI does involve two levels of representation as is advocated here.

2. A formulation with only three kinds of links allows the network to be self-defining to a larger extent, and thus one would expect the corresponding network interpreter to have less "compiled in" knowledge, thereby being both more flexible and less efficient. Also, a formulation with fewer links requires less epistomological committment, whether for better or worse.

3. Since I believe in the use of meta-description, the correspondence of, say, FIDO to the world could be taken as equivalent to the correspondence of the proposition FIDO-EXISTS to the world. This would reduce all correspondence to the correspondence of propositions, but I do not care to operate at the meta-level as this trick requires.

4. We are lead here into the center of the realism/nominalism debate. I embrace the position of Carnap (1950).

5. There has been discussion of whether the referential/attributive ambiguity is a semantic or pragmatic one (e.g. Cole 1978). If we take a semantic distinction to involve two different data structures in the sentence representation and a pragmatic one to involve two different data structures in the discourse, then our solution is a pragmatic one.

6. A method for the sequential binding of quantified variables in determining the truth of an expression was introduced by Hintikka (1976). Although he was working with predicate calculus, he obviously had in mind the same general strategy proposed here.

7. In addition to introducing an entity to the listener, the speaker can indicate whether he expects the listener to learn more about the entity by a choice of <u>this</u>, <u>a</u>, or <u>some</u>. For example,

   Where did you get that balloon?

   a) <u>This man</u> in the park gave it to me.
   b) <u>A man</u> in the park gave it to me.

c) Some man in the park gave it to me.

8. The possibility of making widespread use of the referential/attributive distinction is discussed by Partee (1972).

9. A development of the problem of opaque contexts can be had in Linsky (1971). Quine was unhappy with his (thus our) treatment because it treats intensions (or concepts) as objects. I am not bothered by treating concepts as objects because I believe that any intelligent problem solving system will have to have knowledge about its own knowledge. Besides, circumventing concepts as objects may be computationally as unattractive as retaining them is philosophically. For a recent attempt, which uses reasoning in terms of possible worlds, see Moore (1979). This treatment will face the frame problem (McCarthy and Hayes 1969). For some thoughts of mine, see Martin (1979).

10. The idea, used here, of treating referentially opaque belief contexts as attributive has been proposed by Cole (1978). He points out that the quantifier scope analysis of

John told me that Carol said that the man who killed Kennedy wore size twelve galoshes.

should be three-ways ambiguous, according to whether the quantifier is inside the scope of "told" and "said", just "told", or neither. He suggests that simple application of the referential/attributive distinction would produce only two readings - the attributive reading where "the man who killed Kennedy" is part of Carol's statement and the referential reading which binds "the man who killed Kennedy" in the listener's discourse model. Here, he says, is a discrepancy between the two treatments. One predicts three readings, the other, two.

Cole goes on to suggest that the sentence

John told me, "Carol said that the man who killed Kennedy wore size twelve galoshes."

has only two readings: the attributive reading where "the man who killed Kennedy" is part of Carol's statement and the referential reading where the description "the man who killed Kennedy" is attributed to John. The reading where that description is attributed to the person quoting him is missing. But Cole fails to consider that "the man who killed Kennedy" could be used by Carol referentially.

11. Cole observed that the choice of who to credit with an expression is influenced by the verb; "believe" and "know" favor paraphrase, "say" and "tell" are neutral, "whisper" and "murmur" favor quotation.

12. Partee (1972) points out that this type of sentence and those of the previous section pose a problem for the analysis of the indefinite article in terms of the features ± specific. (Fillmore 1967) Something more is needed to produce all the readings. Features can be used as follows: if a hearer has decided a node is definitely attributive it can be marked +attributive, if it is referential the co-descriptor is shown. Otherwise, it is unmarked and has no co-descriptor.

13. Readers familiar with iteration macros and programming languages will see that Wood's FOR is quite limited, e.g. iteration is limited to a single variable. I use it here because it is simple, is well explained by Woods, and will facilitate comparison of my suggestions with Wood's earlier work.

14. Further on I will discuss both i) the influence of quantifiers in choosing an option and ii) examples with side effects which require consideration of how an iteration could be done.

15. To carry out the iteration the generic BOY would be individuated by a counter and each of these individuals would inherit the description of BOY.

16. Foder (1979a) uses these sentences to argue for the existence of referential readings of indefinites.

17. Also quantifiers can be attached directly to the predicate

    John and Fred <u>each</u> gave Mary $1000.

18. The idea that generic interpretation is a mode of inference is advanced by Nunberg and Pan (1975).

19. This example of Jackendoff (1972) is discussed by Fauconnier (1975).

20. See Webber (1978), Sidner (1979) and Fauconnier (1978) for examples and discussion.

OFFICIAL DISTRIBUTION LIST

Defense Documentation Center
Cameron Station
Alexandria, VA 22314
12 copies

Office of Naval Research
Information Systems Program
Code 437
Arlington, VA 22217
2 copies

Office of Naval Research
Branch Office/Boston
Building 114, Section D
666 Summer Street
Boston, MA 02210
1 copy

Office of Naval Research
Branch Office/Chicago
536 South Clark Street
Chicago, IL 60605
1 copy

Office of Naval Research
Branch Office/Pasadena
1030 East Green Street
Pasadena, CA 91106
1 copy

New York Area
715 Broadway - 5th floor
New York, N. Y. 10003
1 copy

Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D. C. 20375
6 copies

Assistant Chief for Technology
Office of Naval Research
Code 200
Arlington, VA 22217
1 copy

Office of Naval Research
Code 455
Arlington, VA 22217
1 copy

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
(Code RD-1)
Washington, D. C. 20380
1 copy

Office of Naval Research
Code 458
Arlington, VA 22217
1 copy

Naval Ocean Systems Center, Code 91
Headquarters-Computer Sciences &
Simulation Department
San Diego, CA 92152
Mr. Lloyd Z. Maudlin
1 copy

Mr. E. H. Gleissner
Naval Ship Research & Development Center
Computation & Math Department
Bethesda, MD 20084
1 copy

Captain Grace M. Hopper (008)
Naval Data Automation Command
Washington Navy Yard
Building 166
Washington, D. C. 20374
1 copy

Mr. Kin B. Thompson
Technical Director
Information Systems Division
(OP-91T)
Office of Chief of Naval Operations
Washington, D. C. 20350
1 copy

Captain Richard L. Martin, USN
Commanding Officer
USS Francis Marion (LPA-249)
FPO New York, N. Y. 09501
1 copy